GPTune Users Guide[*]

Wissam M. Sid-Lakhdar [†]   Younghyun Cho[‡]   James W. Demmel[‡]

Hengrui Luo[†‡]   Xiaoye S. Li[†]   Yang Liu[†]   Osni Marques[†]

April 13, 2021

Contents

Abstract

GPTune is an autotuning framework that relies on multitask and transfer learning to help solve an underlying black-box optimization problem. GPTune is part of the xSDK4ECP project supported by the Exascale Computing Project (ECP).

Multitask learning and transfer learning have proven to be useful in the field of machine learning when additional knowledge is available to help a prediction task. We aim at deriving methods following these paradigms for use in autotuning, where the goal is to find the best parameters for optimal performance of an application treated as a black-box function.

We assume that evaluations are expensive, and so try to minimize the number of evaluations.

1 Introduction

The goal of autotuning is to automatically choose tuning parameters to optimize the performance of an application. Performance is most often measured by runtime, but any quantitative, measurable quantity is possible, such as the number of messages communicated, amount of memory used, accuracy, etc. Autotuning is particularly challenging when the number of (combinations of) tuning parameter values is large, the performance is a complicated and hard-to-model function of all the tuning parameters, and running the application in order to measure the actual performance is expensive.

GPTune addresses these autotuning challenges by using Bayesian optimization (Gaussian Processes) to build a performance model (by running the application and measuring its performance at a few carefully chosen tuning parameter values), and optimizing the model (eg choosing the tuning parameters to minimize the runtime predicted by the model). In the simplest case, the user has a single task t (eg a linear algebra operation on a matrix of a certain size), the model $f(x)$ predicts the true runtime $y(x)$ as a function of a tuning parameter x (eg a block size), and GPTune chooses x to minimize $f(x)$.

GPTune goes beyond this simple case by providing both Multitask Learning Autotuning (MLA) and Transfer Learning Autotuning (TLA). MLA means using performance data from multiple tasks (eg a fixed linear algebra operation on matrices of several dimensions $t_1, t_2, \ldots, t_k$) to build a more accurate performance model $f(t, x)$ of the true runtime $y(t, x)$ to use for tuning (by choosing x to minimize $f(t_i, x)$, for any t_i). In the common case that the performance varies reasonably smoothly as a function of both t and x, using all the available data to build $f(t, x)$ can make it more accurate, and so better for tuning. Once a predicted optimal x_{opt} is chosen for a particular t_i, the application can be run to measure the actual performance for the values $t = t_i$ and $x = x_{opt}$, and this data is used to update the model $f(t, x)$ to make it more accurate; this process of predicting x_{opt}, measuring the actual performance, and updating the model with the new x sample, can be repeated a user-selected number of times.

TLA goes beyond MLA by using the performance model to predict (and optimize) the performance for tasks (eg values of t) for which no true performance data has been collected. Here we assume that a MLA performance model without using the performance data of these new tasks has been built. The simplest way to do this is to use the same Gaussian Process approach used above to build a model of $x_{opt}(t) = \arg\min_x f(t, x)$, using the known (predicted) values of $\arg\min_x f(t_i, x)$ from MLA. This requires no additional collection of actual performance data. Alternatively, TLA can collect extra performance data for the new tasks, update (rather than rebuild) the MLA model, and predict optimal x for the new tasks (future work).

To simplify the above description, we chose a simple example with one parameter (t) needed to describe the task, and one tuning parameter (x). In practice, there may be many parameters needed to describe the task, and many tuning parameters, which could be of different types (real, integer, or "categorical", i.e. a list of discrete possibilities, such as choices of algorithms). Therefore, our interface (described in detail below) needs to accommodate all these possibilities.

We also need to accommodate constraints on the tuning parameters. For example, the user may require that the number of processors used (a tuning parameter) be less than or equal to an upper bound they supply, or that the product of the number of "processor rows" and the number of "processor columns" be less than the same upper bound. Our interface accommodates these and other possibilities.

If the measured performance data consists of multiple quantities, such as (runtime, accuracy), then the user may want to perform multi-objective optimization. For example, in the case of runtime and accuracy, which are likely to tradeoff against one another (faster runtime leading to worse accuracy), then the user may want to compute the Pareto front of these two quantities. Again, our interface accommodates this and other possibilities.

It is often the case that runtime data will be collected over time by one or more users running many different cases of a popular application, so we want to take advantage of all this data to improve the accuracy of our performance model while reducing the cost of expensive black-box optimization. GPTune provides a history database that allows historical performance data to be stored in files, and reused later. In addition, using a database allows check-pointing, and a more flexible user interface. Moreover, to harness the power of crowd-tuning, we provide a public shared database at `https://gptune.lbl.gov`, where users can store their performance data or download the performance data provided by other users.

To make it easier for users to try different autotuners, since several are available or under active development, our interface allows the user to invoke them as well. So far, OpenTuner [12], HpBandSter [9] and ytopt [19] are supported.

It is sometimes the case that a user has a possibly coarse performance model that can be used to help prune the search space, even if it is not very accurate. Our interface allows the user to submit multiple models, for example just counting arithmetic operations, words communicated, etc., all of which are incorporated in the model GPTune builds, and can help accelerate tuning.

To simplify notation, in the rest of this manual the phrase "task parameter" will refer to an input, like t above, that defines the task to be solved; there are generally multiple task parameters needed to define a task. The phrase "tuning parameter" will refer to a parameter, like x above, that the user wants to optimize; again there are generally multiple parameters to be tuned. The phrase "parameter configuration" will refer to a tuple of a particular setting of the tuning parameters. The word "output" will refer to the performance metric being optimized, such as time.

Table 1 summarizes the notations used in the manual. As an illustrative example, the *QR factorization* routine of *ScaLAPACK* [5], denoted as *PDGEQRF*, is used as an application code to be tuned assuming fixed numbers of compute nodes (*nodes*) and cores per node (*cores*). So we list its respective parameters in this table. Note that only independent task and tuning parameters are listed here. Other parameters, such as the number of threads *nthreads* (used in BLAS) and number of column processes q, can be calculated as $nthreads = \lfloor cores/npernode \rfloor$ and $q = \lfloor nodes * npernode/p \rfloor$. Note that we can use modified parameter definitions to better enforce certain constraints. For example, assuming *cores* is power of 2, we can use a modified parameter $npernode'$ such that $npernode = 2^{npernode'}$.

Symbol	Interpretation
\multicolumn{2}{c}{General notations}	
$\mathbb{IS}$	Task Parameter **I**nput **S**pace
$\mathbb{PS}$	Tuning **P**arameter **S**pace (parameter configurations)
$\mathbb{OS}$	**O**utput **S**pace (e.g., runtime)
$\mathbb{MS}$	performance **M**odel **S**pace (e.g., flop count)
α	dimension of $\mathbb{IS}$
β	dimension of $\mathbb{PS}$
γ	dimension of $\mathbb{OS}$
$\tilde{\gamma}$	dimension of $\mathbb{MS}$
δ (NI)	number of tasks
ϵ (NS, NS1)	number of samples per task
$T \in \mathbb{IS}^{\delta}$	array of tasks selected from sampling
$X \in \mathbb{PS}^{\delta \times \epsilon}$	array of samples (parameter configurations)
$Y \in \mathbb{OS}^{\delta \times \epsilon}$	array of output results (e.g. runtime)
\multicolumn{2}{c}{Example: parameters for ScaLAPACK *PDGEQRF* notations}	
Task m	number of matrix rows
n	number of matrix columns
Tuning mb	row block size
nb	column block size
$npernode$	number of MPI processes per compute node
p	number of row processes

Table 1: Notations. The symbols in the parentheses denote code notations.

At a higher level, we also need to distinguish two independent sets of parameters:

- The parameters associated with the application codes, as described above, i.e. task parameters (as input to GPTune) and tuning parameters (as output from GPTune). Section 4.7 shows the API for the user to define these parameters.

- The parameters associated with the tuner itself, referred to as "GPTune parameters". These are related to the various tuning algorithms, such as MLA and TLA. Sections 4.13, 4.14 and 4.19 show the APIs for the user to choose the GPTune parameters. The settings of these parameters can affect the speed and accuracy of the tuning algorithms.

We have parallelized the most time-consuming part of the GPTune algorithms. Section 3.2 describes the parallel programming model on multicore nodes using MPI and OpenMP. The user can define the parallel computer configuration (nodes, cores) for GPTune's internal computation. The application to be tuned may also be executed in parallel, but GPTune and the objective function may use different numbers of nodes and cores. Section 4.9 shows the API for the user to specify the computer resources.

The rest of this user manual is organized as follows. Section 2 says how to install the system. Section 3 describes the underlying autotuning algorithms and their implementation. Section 4 describes the user interface. Section 5 shows examples of autotuning real applications. Section 6 presents some performance results.

2 Installation

GPTune is implemented in Python and C, and it depends on several Python, Fortran and C packages. Most of them can be installed with one line; the rest requires manual installation. Before the installation, the following software environment is needed (with the required minimum version number in parentheses): **gcc**(7.4.0) or **intel icc**(19.0.0), **openmpi**(4.0.1), **python**(3.7), **scalapack**(2.1.0), **git**, **cmake**(3.19), **blas**, **lapack**. Note that OpenMPI, Python and ScaLAPACK (shared build is recommended) should use the same gcc version. It might be necessary to build them from source. We highly recommend modifying the example build scripts for known machines for the correct installation of GPTune. Alternatively, a pre-built Docker image is also available if the user wants to quickly try the functionality of GPTune.

2.1 Installation using example scripts

GPTune can be obtained from its Github repository:

```
1 git clone https://github.com/gptune/GPTune.git
2 cd GPTune
```

The following example build scripts are available for a collection of tested systems.

- **Ubuntu/Debian-like systems supporting apt-get**. The following script installs everything from scratch and can take up to 2 hours depending on the users' machine specifications. If "MPIFromSource=0", you need to set PATH, LIBRARY_PATH, LD_LIBRARY_PATH and MPI compiler wrappers when prompted.

```
1 bash config_cleanlinux.sh
```

- **Mac OS supporting homebrew.** The following script installs everything from scratch and can take up to 2 hours depending on the users' machine specifications. The user may need to set pythonversion, gccversion, openblasversion, lapackversion on the top of the script to the versions supported by your homebrew software.

```
1 zsh config_macbook.zsh
```

- **NERSC Cori**. The following script installs GPTune with mpi, python, compiler and cmake modules on Cori. Note that you can set "proc=haswell or knl", "mpi=openmpi or craympich" and "compiler=gnu or intel". Setting mpi=craympich will limit certain GPTune features. Particularly, only the so-called reverse communication interface mode can be used, see Section 3.2.2 for more details.

```
1 bash config_cori.sh
```

In all the examples scripts above, setting "BuildExample=0" will only install GPTune with a Scalapack *PDGEQRF* tuning example. Otherwise, all tuning examples (Superlu_DIST, Hypre, STRUMPACK, ButterflyPACK, MFEM) will be installed.

2.2 Installation from scatch

Instead of using the example build scripts, the users can choose to install GPTune step-by-step. The following installs GPTune on a Linux system with gcc compiler, without installing most of the tuning examples.
Obtain GPTune from github.

```
1 git clone https://github.com/gptune/GPTune.git
2 cd GPTune
3 export GPROOT=$PWD
```

Set the following environment variables makes it convenient for consequent installation.

```
1 export PYTHONWARNINGS=ignore
2 export CCC=path-to-the-mpicc-wrapper
3 export CCCPP=path-to-the-mpic++-wrapper
4 export FTN=path-to-the-mpif90-wrapper
5 export RUN=path-to-the-mpirun-wrapper
6 export BLAS_LIB=path-to-the-blas-lib
7 export LAPACK_LIB=path-to-the-lapack-lib
8 export SCALAPACK_LIB=path-to-the-scalapack-lib
```

2.2.1 Python packages in the requirement file

The following python packages listed in requirement.txt can be installed automatically with pip: **numpy, joblib, scikit-learn, scipy, pyaml, matplotlib, GPy, openturns, lhsmdu, pygmo, ipyparallel, opentuner, hpbandster, filelock**.

```
1 pip install --upgrade --user -r requirements.txt
```

2.2.2 GPTune C code

```
1 cd $GPROOT
2 mkdir -p build
3 cd build
4 cmake .. \
5 -DBUILD_SHARED_LIBS=ON \
6 -DCMAKE_CXX_COMPILER=$CCCPP \
7 -DCMAKE_C_COMPILER=$CCC \
8 -DCMAKE_Fortran_COMPILER=$FTN \
9 -DTPL_BLAS_LIBRARIES=$BLAS_LIB\
10 -DTPL_LAPACK_LIBRARIES=$LAPACK_LIB \
11 -DTPL_SCALAPACK_LIBRARIES=$SCALAPACK_LIB
12 make
13 cp lib_gptuneclcm.so ../.
14 cp pdqrdriver ../
```

2.2.3 mpi4py

```
1 cd $GPROOT
2 rm -rf mpi4py
3 git clone https://github.com/mpi4py/mpi4py.git
4 cd mpi4py/
```

```
5 python setup.py build --mpicc="$CCC -shared"
6 python setup.py install
7 export PYTHONPATH=$YTHONPATH:$GPROOT/mpi4py/
```

2.2.4 scikit-optimize

```
1 cd $GPROOT
2 rm -rf scikit-optimize
3 git clone https://github.com/scikit-optimize/scikit-optimize.git
4 cd scikit-optimize/
5 python setup.py build
6 python setup.py install --user
```

2.2.5 autotune

```
1 cd $GPROOT
2 rm -rf autotune
3 git clone https://github.com/ytopt-team/autotune.git
4 cd autotune/
5 cp ../patches/autotune/problem.py autotune/.
6 pip install --user -e .
7 export PYTHONPATH=$YTHONPATH:$GPROOT/autotune/
```

2.2.6 GPTune Examples (SuperLU_DIST)

```
1  cd $GPROOT/examples/SuperLU_DIST
2  git clone https://github.com/xiaoyeli/superlu_dist.git
3  cd superlu_dist
4  mkdir -p build
5  cd build
6  cmake .. \
7  -DCMAKE_CXX_FLAGS="-Ofast -std=c++11 -DAdd_ -DRELEASE" \
8  -DCMAKE_C_FLAGS="-std=c11 -DPRNTlevel=0 -DPROFlevel=0 -DDEBUGlevel=0" \
9  -DBUILD_SHARED_LIBS=OFF \
10 -DCMAKE_CXX_COMPILER=$CCCPP \
11 -DCMAKE_C_COMPILER=$CCC \
12 -DCMAKE_Fortran_COMPILER=$FTN \
13 -DTPL_BLAS_LIBRARIES=$BLAS_LIB \
14 -DTPL_LAPACK_LIBRARIES=$LAPACK_LIB \
15 -DTPL_PARMETIS_INCLUDE_DIRS=path-to-parmetis-include \
16 -DTPL_PARMETIS_LIBRARIES=path-to-parmetis-lib
17 make pddrive_spawn
```

2.3 Docker image

If the users don't want to install GPTune, a Docker image is availabe for all system types: Linux, Mac or Windows. First install and launch docker following instructions at https://docs.docker.com. The docker image can be then obtained and launched by

```
1 docker pull liuyangzhuan/gptune:2.4
2 docker run -it liuyangzhuan/gptune:2.4
```

2.4 Testing the installation

Either of the following two scripts can be tested

```
1 bash run_examples.sh # uncomment corresponding code blocks to activate certain
      tests
2 bash run_ppopp.sh  # edit line 188 accordingly to activiate certain tests (see
      comments inside)
```

Note: if GPTune was installed using config_cleanlinux.sh or the Docker image is loaded, uncomment lines 34-38; if GPTune was installed using config_macbook.zsh, uncomment lines 11-15; if GPTune was installed using config_cori.sh, edit lines 18-22 accordingly.

3 GPTune Implementation

In this section, we describe how GPTune is implemented, with sufficient details of the tuning algorithm so that the user knows how to set GPTune's learning algorithm parameters (see Sections 4.13 and 4.14) and other options (see Section 4.19). We refer the users to our technical paper [18] for a thorough exposition of the methodology.

3.1 Algorithms

We first describe the algorithm for single-objective autotuning, then describe the algorithm for multi-objective autotuning.

3.1.1 Single-objective autotuning

Recall that we seek to optimize some objective function with a set of tunable parameters for best performance, in the form of $\arg\min_x y(t, x)$, where $t \in \mathbb{IS}$ is an input task and $x \in \mathbb{PS}$ is a tuning parameter configuration. $\mathbb{IS}$ is the *Task Parameter Input Space* containing all the input problems that the application may encounter. Note: the word "input" will be dropped in the remaining document. $\mathbb{PS}$ is the *Tuning Parameter Space* containing all the parameter configurations to be optimized, with α being the number of task parameters and β being the number of tuning parameters. We also define $\mathbb{OS}$ to be the *Output Space* of dimension γ, i.e., the number of scalar objective functions.

We can evaluate the objective function pointwise at a tuning parameter configuration (i.e., run and measure the application on a parallel machine), but the function does not have an easy closed form nor easy-to-compute gradients. Moreover it is expensive to perform a function evaluation. Given these characteristics, we choose to use the Bayesian optimization method (also known as Efficient Global Optimization (EGO) [13]), in which a prior belief model representing the assumptions on the objective function is chosen, and a posterior is built from it so as to maximize the likelihood of some probability distribution function based on the sampled objective function values.

Compared to the other autotuning efforts, one of our innovations is to use *multitask learning* to build a more accurate predictive model. Multitask learning consists of learning several tasks simultaneously (eg running a ScaLAPACK routine with different matrix dimensions) while sharing common knowledge between them in order to improve the prediction accuracy of each task and/or speed up the training process. We call this framework multitask learning autotuning (MLA).

Specifically, the MLA learning process consists of the following phases, where we also define GPTune parameters, or provide names for other quantities that appear in the GPTune interface.

1. **Sampling phase.** There are two sampling steps. The first is to select a set T of δ tasks $T = [t_1; t_2; \ldots; t_\delta] \in \mathbb{IS}^\delta$. The goal is to get a representative sample of the variety of problems that the application may encounter, rather than focusing on a specific type of problem. Alternatively, T can represent a list of target tasks specified by the user, instead of sampling done by GPTune.

   The second sampling step is to select an initial set $X = [X_1; X_2; \ldots; X_\delta]$ of tuning parameter configurations for every task. Let NS denote a prescribed total number of function evaluations per task. The number of initial samples is set to $\epsilon = \text{NS1}$, with $\text{NS1} = \lceil \text{NS}/2 \rceil$ by default. For task t_i, its initial sampling X_i consists of ϵ tuning parameter configurations $X_i = [x_{i,j}]_{j \in [1,\epsilon]} \in \mathbb{PS}^\epsilon$. Define $X = [X_1; X_2; \ldots; X_\delta] \in \mathbb{PS}^{\delta \times \epsilon}$ to represent all the samples.

   The samples $x_{i,j}$ are evaluated through runs of the application, whose results, $y_{i,j} = y(t_i, x_{i,j}) \in \mathbb{OS}$, can be formed as $Y_i = [y_{i,j}]_{j \in [1,\epsilon]} \in \mathbb{OS}^\epsilon$. The set $Y = [Y_1; Y_2; \ldots; Y_\delta] \in \mathbb{OS}^{\delta \times \epsilon}$ represents the results of all these evaluations.

   GPTune parameters. NI: number of tasks (i.e., δ). Note that we use NI as the code notation and δ as the algorithm notation. Igiven: (optional) list of user specified task parameters. NS: total number of function evaluations per task. NS1: (optional, default value given above) number of initial function evaluations. The other related GPTune parameters are: sample_class, sample_algo, sample_max1_iter, which are described in Section 4.19.

2. **Modeling phase.** This phase builds a Bayesian posterior probability distribution of the objective function via training a model of the black-box objective function relative to the tasks in T. We derive a single model that incorporates all the tasks, sharing the knowledge between them to be able to better predict them all. To this end, we use the *Linear Coregionalization Model* (LCM), which is a generalization of *Gaussian Process* (GP) in the multi-output setting. A GP represents a probability model $f(x)$ for the objective function $y(x)$. It assumes that $(f(x_1), \ldots, f(x_\epsilon))$ is jointly Gaussian, with mean function $\mu(x)$ and covariance $\Sigma(x, x') = k(x, x')$, where k is a positive definite kernel function. The idea is that if x and x' are deemed similar by the kernel, we expect the outputs of the function at those points to be similar too. A model $f(x)$ following a GP is written as: $f(x) \sim GP(\mu(x), \Sigma)$. In practice, $\mu(x)$ is initialized to be the zero function. The modeling is done through $\Sigma(X, X)$, by maximizing the log-likelihood of the samples X with values Y on the GP. For single task learning, the size of the covariance matrix $\Sigma(X, X)$ is $\epsilon \times \epsilon$.

   The key to LCM is the construction of an approximation of the covariance between the different outputs of the model of every $t_i \in T$. In this method, the relations between outputs are expressed as linear combinations of independent *latent random functions*

$$f(t_i, x) = \sum_{q=1}^{Q} a_{i,q} u_q(x) \tag{1}$$

   where $a_{i,q}$ ($i \in [1, \delta]$) are hyperparameters to be learned, and u_q are the latent functions, each of which is an independent GP whose hyperparameters need to be learned as well.

Due to the independence of u_q's, the covariance between two outputs is simply the sum of auto-covariances of u_q at those two points:

$$cov(f(t_i, x), f(t_{i'}, x')) = \sum_{q=1}^{Q} a_{i,q} a_{i',q} cov(u_q(x), u_q(x')) \tag{2}$$

In LCM, we assume the covariance of the latent function is based on a Gaussian kernel:

$$cov(u_q(x), u_q(x')) = k_q(x, x') = \sigma_q^2 \exp\left(-\sum_{i=1}^{\beta} \frac{(x_i - x_i')^2}{l_i^q} \right) \tag{3}$$

When considering all the tasks and all the samples together, the covariance matrix $\Sigma(X, X)$ is of size $\delta \cdot \epsilon$ with entries

$$\Sigma(x_{i,j}, x_{i',j'}) = \sum_{q=1}^{Q} (a_{i,q} a_{i',q} + b_{i,q} \delta_{i,i'}) k_q(x_{i,j}, x_{i',j'}) + d_i \delta_{i,i'} \delta_{j,j'} \tag{4}$$

where $\delta_{i,j}$ is the Kronecker delta function, $b_{i,q}$ and d_i are diagonal regularization parameters. The learning task in the nth MLA iteration is to find the best hyperparameters of the model, such as the hyperparameters σ_q, l_i^q in the Gaussian kernel Eq.(3) and coefficients $a_{i,q}, d_i$ in Eq. (4). We use a gradient-based optimization algorithm to maximize the log-likelihood of the model on the data. Specifically, we employ the limited-memory Broyden–Fletcher–Goldfarb–Shanno algorithm (L-BFGS) [17] through the Python package scikit-Optimize [3]. Note that the log-likelihood function is usually highly nonconvex, so local optimization may not converge to the/a global optimum. The modeling phase extends the Python package GPy [11] to enable distributed-memory parallel modeling.

GPTune parameters. model_latent: number of latent functions (i.e., Q). By default $Q = \delta$. The other relevant options are model_restarts, model_max_iters, etc. See Section 4.19 for details.

Algorithm 1 Bayesian optimization-based single-objective MLA

1: **Sampling phase:** Compute $y(t_i, x)$, $i \leq \delta$ at NS1 initial random tuning parameter configurations for δ selected tasks. Set $\epsilon = $ NS1.
2: **while** $\epsilon \leq$ NS **do**
3:     **Modeling phase:** Update the hyperparameters in the LCM model of $y(t_i, x)$, $i \leq \delta$ using all available data.
4:     **Search phase:** Search for an optimizer x_i^* for the EI of task t_i, $i \leq \delta$. Let $X^* = [x_1^*; x_2^*; \ldots; x_\delta^*]$.
5:     Compute $y(t_i, x)$, $i \leq \delta$ at the new tuning parameter configurations X^*.
6:     $\epsilon \leftarrow \epsilon + 1$.
7: **end while**
8: Return the optimum tuning parameter configurations and objective function values for each task.

---

3. **Search phase.** Once the model has been updated, the objective function values at new points $X^* = [x_1^*; x_2^*; \ldots; x_\delta^*]$ can be predicted with posterior mean $\mu_* = [\mu_1^*; \mu_2^*; \ldots; \mu_\delta^*]$ and posterior variance (prediction confidence) $\sigma_*^2 = [\sigma_1^{*2}; \sigma_2^{*2}; \ldots; \sigma_\delta^{*2}]$ as:

$$\mu_* = \Sigma(X^*, X)\Sigma(X, X)^{-1}Y \tag{5}$$

$$\sigma_*^2 = \Sigma(X^*, X^*) - \Sigma(X^*, X)\Sigma(X, X)^{-1}\Sigma(X^*, X)^T \tag{6}$$

Note that the posterior variance is equal to the prior covariance minus a term that corresponds to the variance removed by observing X [10]. The mean and variance can be used to construct the Expected Improvement (EI) acquisition function, which can be maximized in order to choose a new point X^* for function evaluation (additional sampling). We use evolutionary algorithms provided by the python package PyGMO [2] to optimize the EI. PyGMO supports many optimization algorithms. By default, we use Particle Swarm Optimization (PSO) algorithm [15].

With one additional function evaluation, we increment ϵ by 1 and move to next MLA iteration for model improvement until ϵ reaches a prescribed sample count NS (i.e., a prescribed budget of function evaluations). This iterative process is summarized as Algorithm 1.

GPTune parameters. search_algo: evolutionary algorithms supported by PyGMO. By default search_algo is 'pso'. The other relevant options are search_pop_size, search_gen, search_evolve, search_max_iters, etc. See Section 4.19 for details.

3.1.2   Multi-objective autotuning

The MLA algorithm described in Section 3.1.1 can be easily extended to multi-objective, multi-task settings. Algorithm 2 describes the multi-objective extension of Algorithm 1. Let $y^s(t, x)$, $s \leq \gamma$ denote the sth objective function. Algorithm 2 essentially builds one LCM model per objective function $y^s(t, x)$ in the modeling phase. In addition, the search phase relies on multi-objective evolutionary algorithms such as non-dominated sorting generic algorithm II (NSGA-II) [7] to search for k new tuning parameter configurations in each iteration. The sorting is based on the Pareto dominance and Crowding distance [7].

Algorithm 2 Bayesian optimization-based multi-objective MLA

1: **Sampling phase:** Compute $y^s(t_i, x)$, $i \leq \delta$, $s \leq \gamma$ at NS1 initial random tuning parameter configurations for δ selected tasks. Set $\epsilon = $ NS1.
2: **while** $\epsilon \leq$ NS **do**
3:     **Modeling phase:** For each objective $s \leq \gamma$, update the hyperparameters in the LCM model of $y^s(t_i, x)$, $i \leq \delta$ using all available data.
4:     **Search phase:** Search for k best tuning parameter configurations for the EI of task t_i, $i \leq \delta$.
5:     Compute $y^s(t_i, x)$, $i \leq \delta$ at the k new tuning parameter configurations.
6:     $\epsilon \leftarrow \epsilon + k$.
7: **end while**
8: Return the optimum tuning parameter configurations and objective function values for each task.

GPTune parameters. search_algo: evolutionary algorithms supported by PyGMO. By default search_algo is 'nsga2' (i.e., the NSGA-II algorithm). search_more_samples: the number of additional samples per iteration (i.e., k). The other relevant options are search_pop_size, search_gen, search_evolve, search_max_iters, etc. See Section 4.19 for details.

3.1.3 Incorporation of performance models

A performance model refers to an analytical formula or inexpensive application run for any feature (time, memory, communication volume, flop counts) of the objective function. For example, one can provide an analytical formula for the flop count when the objective function is the runtime. When available, performance models can be incorporated to build a more accurate LCM model with fewer samples needed. In what follows, the performance model incorporation is explained assuming single task $\delta = 1$ and single objective $\gamma = 1$ for simplicity.

We define $\mathbb{MS}$ to be the *performance Model Space* with dimension $\tilde{\gamma}$ being the number of models. Let $\tilde{y}(x)$ denote the results of the performance models for parameter configuration x. Without the performance model, entries of the LCM kernel matrix represents the nonlinear inner products between points x and x' in the feature space $\mathbb{PS}$ of dimension β. One can use the values $\tilde{y}(x)$ as the extra features to construct an enriched feature space of dimension $\beta + \tilde{\gamma}$ consisting points $[x, \tilde{y}(x)]$. Note that the enriched LCM matrix still has the same dimension $\epsilon\delta \times \epsilon\delta$. Once the LCM model is built, the objective function at the new point x^* can still be predicted using (5) and (6) by replacing x^* with $[x^*, \tilde{y}(x^*)]$.

Note that GPTune internally represents the tuning parameters x as real numbers in the unit hypercube $[0, 1]^{\beta}$. Therefore, it is beneficial that the $\tilde{y}(x)$ can be normalized by the users to $[0, 1]^{\tilde{\gamma}}$. For example, if the users know in advance that $\tilde{y}(x) \in [y_l, y_u]$ assuming $\tilde{\gamma} = 1$, they can instead return the output of the model as $\tilde{y}(x) \to (\tilde{y}(x) - y_l)/(y_u - y_l)$. We recommend the users to perform such conversion even when y_l and y_u are approximate numbers, as this will help GPTune build a more accurate LCM model.

3.1.4 Transfer learning

Transfer learning autotuning (TLA) focuses on the case where no true performance data has been collected for a specified task t^*, but rather performance data and models have been built for different (but related) tasks $t_i, i \leq \delta$. The simplest way to do this is to use the Gaussian Process approach to build a model of $x_{opt}(t) = \arg\min_x f(t, x)$, using the known (predicted) values of $\arg\min_x f(t_i, x)$ from MLA. This requires no additional collection of actual performance data. GPTune uses Algorithm 1 with $\delta = 1$ to build this model.

Alternatively, the model could be iteratively used to predict an optimal x, measure the performance, and update the model. This is future work.

3.2 Parallel implementations

GPTune supports both shared-memory and distributed-memory parallelism through dynamic thread and process management. Most parts of GPTune are implemented with Python3. The shared-memory parallelism is supported through OpenMP threading or the subclass ThreadPoolExecutor from the Python module *concurrent.futures*. The distributed-memory parallelism is supported through the Python package mpi4py [1] and will be explained in detail.

3.2.1 Dynamic process management

In our design, only one MPI process can execute the GPTune driver, but it can also dynamically create new groups of MPI processes (workers) to speed up the objective function evaluation, modeling phase and search phase through the use of MPI spawning. To describe the spawning mechanism, we recall that there are two kinds of MPI communicators, i.e. intra- and inter-communicators. An intra-communicator consists of a group of processes and a communication context, while an inter-communicator binds a communication context with two groups (local and remote) of processes. The master process (running the GPTune driver) will call the function "Spawn" in mpi4py to create a group of new processes. The master process is contained in the intra-communicator "MPI_World" with only one process. The Spawn function will return an inter-communicator "SpawnedComm" that contains a local group (the master itself) and a remote group (containing the workers). The workers also have their own intra-communicator "MPI_World" and call the mpi4py function "Get_parent" that returns an inter-communicator "ParentComm" that contains a local group (the workers) and a remote group (the master). Data can be communicated between the master and workers using the inter-communicators. This scheme can be conceptually depicted in Fig. 1.

Figure 1: GPTune parallel programming model.

A typical spawning call on the master process is

```
1  SpawnedComm = mpi4py.MPI.COMM_SELF.Spawn(executable, args, maxprocs,info)
```

Here, "maxprocs" is the number of new processes to be created, "executable" is the program to be executed by the workers, "args" are the command line arguments to be passed to the program, "info" are the environment variables to be passed to the program, and "SpawnedComm" is the inter-communicator with the master as the local group.

A typical setup on the worker processes is

```
1  ParentComm = mpi4py.MPI.Comm.Get_parent()
```

Here "ParentComm" is the inter-communicator with the workers as the local group. Note that Get_parent can be called from C, C++ or Fortran application codes with a slightly different syntax.

In what follows, we describe the shared-memory and distributed-memory parallelism in the objective function evaluation, modeling phase (in MLA), and search phase (in MLA) separately. Parallel implementation of TLA is future work.
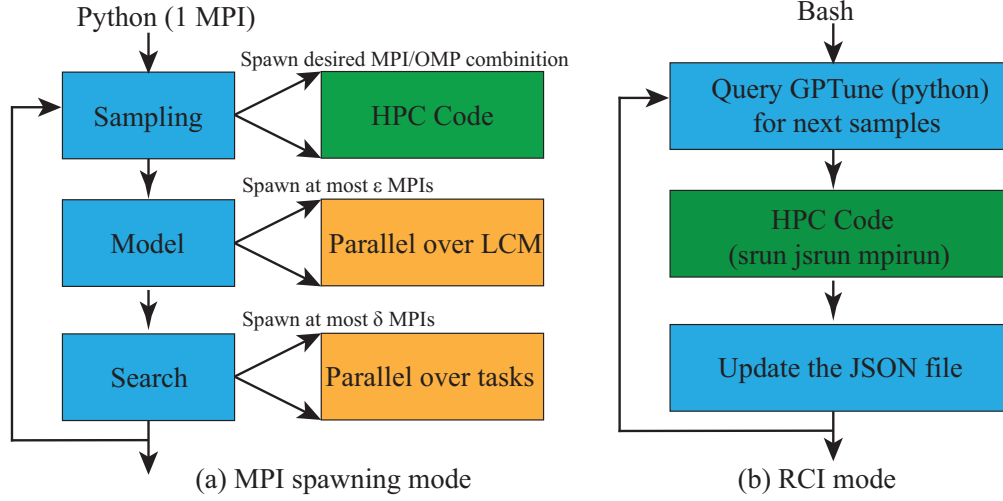
Figure 2: Two execution modes of GPTune: (a) MPI spawning-based interface (the default) and (b) reverse communication interface (RCI)

3.2.2 Objective function evaluation

GPTune provides two interfaces for objective function evaluation, i.e., the MPI spawning mode (the default) and the reverse communication interface (RCI) mode.

For the default mode, the (MPI) application code needs to be compiled using the same software dependence as GPTune (OpenMPI, scalapack, etc.), and with the insertion of a few extra lines (see Section 4.1). This makes the application code callable from inside the GPTune Python modules (see Figure 2 left).

On the other hand, the RCI mode doesn't require OpenMPI-compiled application code and requires a bash script to invoke GPTune and the application. More specifically, the bash script will query GPTune (a Python driver) for next sample points, search for required samples in the database, invoke the application code, write the evaluation results into the database, and then call the GPTune python driver again to ask for next samples (see Figure 2 right). As such, the parallelism of the objective function evaluation is completely handled in the bash script. The RCI mode will be further explained in Section 4.18 and 5.3. The rest of Section 3.2.2 only applies to the default mode.

In the default mode, the users of GPTune will need to provide the objective function (see the definition in Section 4.2) that uses the task and tuning parameters to execute the application code. There are two levels of parallelism supported: running a single objective function evaluation with given MPI and thread counts, and running multiple objective function evaluations in parallel.

Running one function evaluation in parallel. For a single parallel objective function evaluation, the MPI count can be passed to the application code using the argument "maxprocs", the thread count can be passed using the argument "info" as

```
info = mpi4py.MPI.Info.Create()
info.Set('env', 'OMP_NUM_THREADS=%d\n' %(nthreads))
```

where "nthreads" is the number of threads possibly contained in the task or tuning parameters. Depending on how the application code is implemented, one can pass the parameters using command line ("args"), environment variables ("info"), or an input file stored on disc.

To collect the returning value(s) from the workers, one can choose to read from the individual log file per function execution (see the example of ScaLAPACK QR in Section 5.1) or communicate using the inter-communicators (see the example of SuperLU_Dist in Section 5.2). For example, assuming the application code is written in C, the two suggested ways of communicating data are:

Using log file: On the master side, write the (Python) objective function (see Section 4.1) in the following fashion:

```
1 def objectives(point):
2 # extract task and tuning parameters from "point", and use them to define
      enviroment variables (info), command line arguments (args), MPI counts (
      maxprocs) and input files if needed
3 info = mpi4py.MPI.Info.Create() # enviroment variables
4 envstr= 'env1=%d\n' %(ev1)
5 envstr+= 'env2=%d\n' %(ev2)
6 info.Set('env',envstr)
7 args = ['-a1', '%s'%(av1), '-a2', '%s'%(av2)] # command line arguments
8 create the input file needed by executable  # input file
9 SpawnedComm = mpi4py.MPI.COMM_SELF.Spawn(executable, args, maxprocs,info) # call
      executable
10 SpawnedComm.Disconnect() # destroy inter-communicator
11 read objective values from the log file into res  # output file
12 return res
```

On the worker side, the C function looks like the following:

```
1 int main(int argc, char *argv[]){
2 MPI_Init( &argc, &argv );
3 MPI_Comm parent; MPI_Comm_get_parent(&parent); /* Get the inter communicator. */
4 .../* Read the parameters from environment variable, command lines and/or input
      files, compute the objective function and dump it into a log file. */
5 MPI_Comm_disconnect(&parent); /* Disconnect the inter communicator. */
6 MPI_Finalize();}
```

Using inter-communicator: For example, the master can collect data using MPI_Reduce as

```
1 SpawnedComm.Reduce(sendbuf=None,recvbuf=data,op=MPI_OP,root=mpi4py.MPI.ROOT)
```

and the workers send the data by

```
1 ParentComm.Reduce(sendbuf=data,recvbuf=None,op=MPI_OP,root=0)
```

Here, "data" is the returning value, "MPI_OP" is the MPI reduce operation. Again, the syntax can be different on the workers depending on the programming language of the application code. One can also consider using MPI_Send and MPI_Recv for passing the data back to the master, please refer to the mpi4py documentation for more details. One can consider modifying the following example: On the master side, the (Python) objective function (see Section 4.1) looks like the following:

```
1 def objectives(point):
2 # extract task and tuning parameters from "point", and use them to define
      enviroment variables (info), command line arguments (args), MPI counts (
      maxprocs) and input files if needed
3 info = mpi4py.MPI.Info.Create() # enviroment variables
```

```
4  envstr= 'env1=%d\n' %(ev1)
5  envstr+= 'env2=%d\n' %(ev2)
6  info.Set('env',envstr)
7  args = ['-a1', '%s'%(av1), '-a2', '%s'%(av2)] # command line arguments
8  create the input file needed by executable  # input file
9  SpawnedComm = mpi4py.MPI.COMM_SELF.Spawn(executable, args, maxprocs,info) # call
       executable
10 SpawnedComm.Reduce(sendbuf=None,recvbuf=data,op=MPI_OP,root=mpi4py.MPI.ROOT) # use
        MPI_Reduce for collect the objectives
11 SpawnedComm.Disconnect() # destroy inter-communicator
12 return res
```

On the workder side, the C function looks like the following:

```
1  int main(int argc, char *argv[]){
2  MPI_Init( &argc, &argv );
3  MPI_Comm parent; MPI_Comm_get_parent(&parent); /* Get the inter communicator. */
4  .../* Read the parameters from environment variable, command lines and/or input
       files, and compute the objective function. */
5  MPI_Reduce(result, MPI_BOTTOM, 1, MPI_FLOAT,MPI_MAX, 0, parent); # return the
       values "result" to the master
6  MPI_Comm_disconnect(&parent); /* Disconnect the inter communicator. */
7  MPI_Finalize();}
```

Running multiple function evaluations in parallel   For applications that require a small to modest core count for each objective function evaluation, it's beneficial to even perform multiple function evaluations simultaneously. GPTune uses MPI spawning and ThreadPoolExecutor to support distributed- and shared-memory parallelism over number of function evaluations. Note that the users need to use this feature with caution: one needs to make sure the output and input files (when they exist) from different function evaluations do not interfere with each other. This can be done by e.g., assigning each function evaluation a unique file or directory name.

GPTune parameters. objective_nprocmax: maximum core counts for each function evaluation. objective_evaluation_parallelism: whether to use parallelism over different function evaluations. objective_multisample_processes: when both objective_evaluation_parallelism and distributed_memory_parallelism are set to True, this parameter denotes the number of process groups, each responsible for one function evaluation. objective_multisample_threads: when both objective_evaluation_parallelism and shared_memory_parallelism are set to True, this parameter denotes the number of threads processes for parallelizing over the number of function evaluations. See Section 4.19 for details.

3.2.3 Modeling phase of MLA

The modeling phase, described in Section 3.1.1, uses the L-BFGS algorithm to find a set of LCM hyperparameters that minimizes the log-likelihood function using selected objective function samples. The GPTune implementation can choose n_{start} random starting guesses of the hyperparameters, each used by L-BFGS to search for the minimum log-likelihood. GPTune then chooses the set of hyperparameters that yields the best log-likelihood and finishes the modeling phase.

The current implementation supports two levels of parallelism in this phase: (1) The number of n_{start} random starts and corresponding L-BFGS optimization are distributed over user specified

17

number of threads or MPI processes. Note that the shared-memory parallelism and distributed-memory parallelism (over n_{start}) are supported mutually exclusively. GPTune uses MPI spawning to support the distributed-memory parallelism for the random starts. (2) For each L-BFGS optimization, the factorization of the covariance matrix is parallelized over user specified number of threads and MPI processes, the formation of the covariance matrix is parallelized over user specified number of threads. GPTune uses MPI spawning for distributed-memory parallelization of the covariance matrix.

GPTune parameters. model_restarts: number of random starts of the hyperparameters (i.e., n_{start}). distributed_memory_parallelism: whether to use distributed-memory parallelism over the random starts. model_restart_processes: number of MPI processes for parallelizing over n_{start}. shared_memory_parallelism: whether to use shared-memory parallelism over the random starts. model_restart_threads: number of threads for parallelizing over n_{start}. model_processes: number of MPI processes for parallelizing factorization of the covariance matrix. model_threads: number of OpenMP threads for parallelizing formation and factorization of the covariance matrix. See Section 4.19 for details.

3.2.4   Search phase of MLA

The search phase uses evolutionary algorithms in PyGMO to search for the next sample point in each task (see Section 3.1.1 for details). The GPTune implementation supports two levels of parallelism: (1) The multi-task search can be parallelized over the δ tasks using user-specified number of threads or MPI processes. Note that the shared-memory parallelism and distributed-memory parallelism (over δ) are supported mutually exclusively. GPTune uses MPI spawning to support the distributed-memory parallelism over the tasks. (2) For each task, the evolutionary algorithm can be parallelized using user-specified number of threads.

GPTune parameters. distributed_memory_parallelism: whether to use distributed-memory parallelism over the δ tasks. search_multitask_processes: number of MPI processes for parallelizing over δ. shared_memory_parallelism: whether to use shared-memory parallelism over δ. search_multitask_threads: number of threads for parallelizing over δ. search_threads: number of threads used in PyGMO. See Section 4.19 for details.

3.3   History Database

3.3.1   Design

The GPTune interface (Section 4) allows GPTune to automatically store and load historical performance data to and from performance data files with JavaScript Object Notation (JSON [14]) format in the user's local storage. Each application (tuning problem) has a separate data file that contains all the historical performance data of the tuning problem. Figure 3 illustrates how the history database runs along with GPTune's Bayesian optimization.

Each JSON file contains all function evaluation results obtained from the GPTune's Bayesian optimization model. After evaluating each parameter configuration, GPTune stores the task parameter, the tuning parameter, and the evaluated result into the JSON file. This practice ensures that no data is lost, in the cases where (a) a long run with many parameter configurations does not complete due to time limitation, or (b) certain parameter configuration crashes and causes the tuner to stop. If GPTune is run in parallel and multiple processes attempt to update the JSON file simultaneously, the history database allows only one process to update the file at a time based
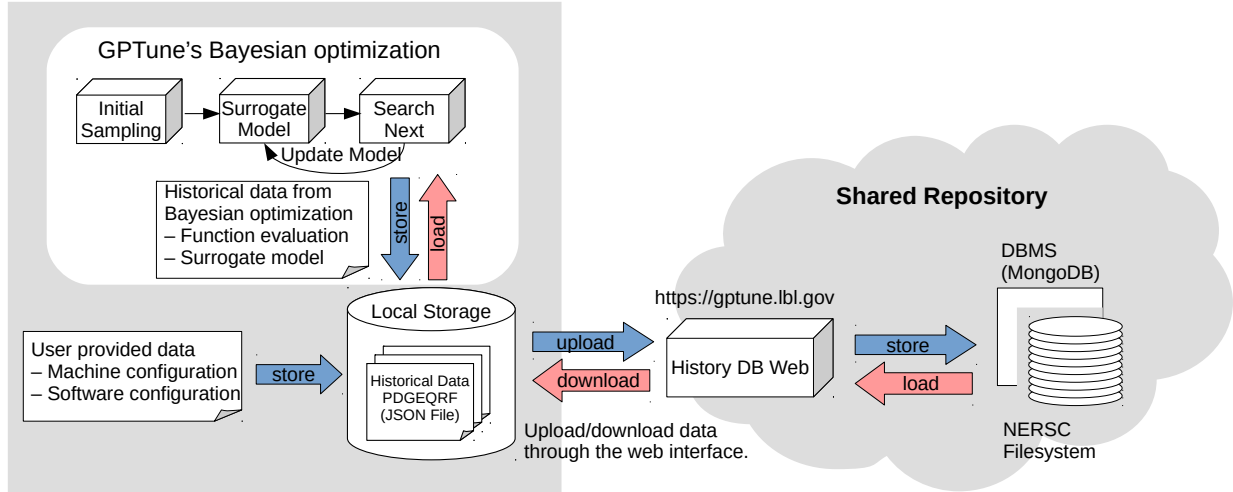
Figure 3: History database design.

on simple file access control (either using Python's filelock module or using `rsync`-based file synchronization). GPTune also records user-provided meta-description like machine configuration and software information (e.g. which software libraries are used for that application) along with the function evaluation results. Based on the meta information, users can determine which data are relevant for learning from a possibly different machine or software versions or configurations. In Section 4, we introduce several GPTune interfaces to provide the software/machine configuration. GPTune also supports storing and loading trained GP surrogate models along with some model statistics information such as likelihood values of the model. Users can use this feature to not only re-use pre-trained models for autotuning, but also analyze model for research purposes (e.g. sensitivity analysis on the tuning parameters).

To harness the power of crowd-tuning, we provide a public shared repository at `https://gptune.lbl.gov` through NERSC (`https://www.nersc.gov`)'s Science Gateways, where users can upload their performance data obtained from GPTune and download performance data provided by other users. In the shared repository, all submitted performance data is stored in a storage provided by NERSC and internally managed by using MongoDB [6]. The shared database requires login credentials for users to submit their performance data. Every submitted performance data can have multiple accessibility options: publicly available data, private data, or data that can be shared with specific users. In other words, the shared repository allows anyone to browse and download publicly available data.

For more details about the history database methodology and the shared repository, we refer the interested users to our online manual at `https://gptune.lbl.gov/docs`.

3.3.2 JSON Format

In this section, we explain the JSON format to store performance data from GPTune. Each tuning problem has a separate data file (e.g. `tuning_problem_name.json`) that contains all performance data (obtained by the user and/or downloaded from the shared public database) of the tuning problem. Each JSON file has two labels `func_eval` and `model_data`. As the name indicates,

19

`func_eval` contains the list of all function evaluation results, and `surrogate_model` contains the list of each trained surrogate model's meta-data.

```json
{
  "func_eval": [
    {
      /* function evaluation result */
    },
    {
      /* function evaluation result */
    },
    ...
  ],
  "surrogate_model": [
    {
      /* surroagte model meta-data */
    },
    {
      /* surroagte model meta-data */
    },

    ...
  ]
}
```

Function Evaluation Result Listing 1 shows a function evaluation result of the *PDGEQRF* routine of ScaLAPACK [5] for a given task/parameter configuration. In the listing, `task_parameter` contains the information about the task parameter, and `tuning_parameter` contains the tuning parameter configuration, and its evaluation result is stored in `output`.

Each function evaluation data can also contain the information about the machine and software configuration to run the application (or the tuning problem). The information related to the machine configuration includes the machine name (e.g. Cori) and the number of cores/nodes used. The software configuration includes the versions of the software packages The software information contains the versions of software packages that are used for compiling/installing the application. The machine and software configurations are stored in `machine_configuration` and `software_configuration`, respectively.

Also, when saving a function evaluation result, the data-creation time and a unique ID of the function evaluation are automatically generated and appended by GPTune. This information can be useful if the user uploads the data into our shared repository. If different users submit function evaluation results for the same task and parameter configurations, the database can differentiate between different function evaluation results based on their UIDs.

```json
{
  "task_parameter": {
    "m": 30000,
    "n": 30000
  },
  "tuning_parameter": {
    "mb": 5,
    "nb": 13,
    "npernode": 2,
```

```
10       "p": 9
11     },
12     "evaluation_result": {
13       "r": 15.148637
14     },
15     "machine_configuration": {
16       "machine": "cori",
17       "haswell": {
18         "nodes": 8,
19         "cores": 32
20       },
21       "knl": {
22         "nodes": 0,
23         "cores": 0
24       }
25     },
26     "software_configuration": {
27       "openmpi": {
28         "version_str": "4.0.0",
29         "version_split": [
30           4,
31           0,
32           0
33         ],
34         "tags": "lib,mpi,openmpi"
35       },
36       "scalapack": {
37         "version_str": "2.1.0",
38         "version_split": [
39           2,
40           1,
41           0
42         ],
43         "tags": "lib,scalapack"
44       }
45     },
46     "time": {
47       "tm_year": 2021,
48       "tm_mon": 1,
49       "tm_mday": 27,
50       "tm_hour": 21,
51       "tm_min": 22,
52       "tm_sec": 22,
53       "tm_wday": 2,
54       "tm_yday": 27,
55       "tm_isdst": 0
56     },
57     "uid": "cc7c03ec-6128-11eb-a40f-85c4081a47e2"
58 }
```

Listing 1: Example Function Evaluation Result

Surrogate Model Listing 2 shows the information of a surrogate model for the IJ routine of Hypre [8] for five different function evaluation results for task {i: 200, j: 200, k: 200}.

Label `hyperparameters` contains the hyperparameters values which are required to reproduce

the surrogate model. Here, we assume the GPTune's default modeling scheme (Section 3.1), based on the LCM, is used. Recall that in LCM, eq (3) and (4), the hyperparameters are l_j^q, $a_{i,q}$, σ_q, $b_{i,q}$, d_i. Hence, in the below example which considers one task ($Q = 1$) for 12 tuning parameters ($\beta = 12$), we store 16 hyperparameters in total (12 for l_j^q and 1 for each of $a_{i,q}$, σ_q, $b_{i,q}$, d_i). As another example, considering two tasks ($Q = 2$) for 12 tuning parameters ($\beta = 12$), we need to store 36 hyperparameters in total (24 for l_j^q, 4 for $a_{i,q}$, 2 for σ_q, 4 for $b_{i,q}$, 2 for d_i).

model_stats stores the model's statistics information. For the GPTune's LCM, we can store some statistics information such as *log_likelihood*, *neg_log_likelihood*, *gradients*, and *iteration* (how many iterations were required to converge the model). Note that, trained surrogate models may or may not be meaningful for different problem spaces. Therefore, the JSON data also contains task parameter information (task_parameters) and which function evaluation results were used (func_eval) to build the surrogate model, by containing the list of the UIDs of the function evaluation results. The history database can load trained models only if they match the problem space of the given optimization problem. Similar to function evaluation results, the data generation time and a unique ID of each surrogate model are also automatically appended by GPTune.

```
1  {
2    "hyperparameters": [
3      0.5140214197473143,
4      1.037247070366763,
5      337.6636254330382,
6      15.04072992869631,
7      5.107732628800839,
8      2.9558180040483086,
9      8.696644421366367,
10     26.085771260561174,
11     25.125605700207174,
12     4.608683353484095,
13     3.511325043265344,
14     1.994878529737146,
15     29.597204778514428,
16     2.3554709171760972,
17     9.999999586881094e-06,
18     69.65076357886942
19   ],
20   "model_stats": {
21     "log_likelihood": -35.22869969421778,
22     "neg_log_likelihood": 35.22869969421778,
23     "gradients": [
24       -0.5379014374164104,
25       0.02347984513751207,
26       3.1845619569323703e-06,
27       0.0019392480264629336,
28       0.022834797398354992,
29       -0.03018478674355551,
30       0.022739720419626193,
31       -0.00010746831896165282,
32       0.002468829248110913,
33       0.01631023320692664,
34       0.030788540386404495,
35       -0.024631175060550625,
36       -0.27274355043218274,
37       -6.198242202047159e-05,
```

```
38          -3.262995507641375e-05,
39          -0.2253219864474227
40        ],
41        "iteration": 65
42      },
43      "func_eval": [
44        "09aab368-612d-11eb-8bf3-bbda784b918d",
45        "09aae608-612d-11eb-8bf3-bbda784b918d",
46        "09ab0c32-612d-11eb-8bf3-bbda784b918d",
47        "09ab30fe-612d-11eb-8bf3-bbda784b918d",
48        "09ab54ee-612d-11eb-8bf3-bbda784b918d"
49      ],
50      "task_parameters": [
51        [
52          200,
53          200,
54          200
55        ]
56      ],
57      "problem_space": {
58        "IS": [
59          {
60            "type": "int",
61            "lower_bound": 20,
62            "upper_bound": 1024
63          },
64          {
65            "type": "int",
66            "lower_bound": 20,
67            "upper_bound": 1024
68          },
69          {
70            "type": "int",
71            "lower_bound": 20,
72            "upper_bound": 1024
73          }
74        ],
75        "PS": [
76          {
77            "type": "int",
78            "lower_bound": 1,
79            "upper_bound": 31
80          },
81          {
82            "type": "int",
83            "lower_bound": 1,
84            "upper_bound": 31
85          },
86          {
87            "type": "int",
88            "lower_bound": 30,
89            "upper_bound": 31
90          },
91          {
92            "type": "real",
93            "lower_bound": 0,
```

```
 94          "upper_bound": 1
 95        },
 96        {
 97          "type": "real",
 98          "lower_bound": 0,
 99          "upper_bound": 1
100        },
101        {
102          "type": "int",
103          "lower_bound": 1,
104          "upper_bound": 12
105        },
106        {
107          "type": "categorical",
108          "categories": [
109            "0",
110            "1",
111            "2",
112            "3",
113            "4",
114            "6",
115            "8",
116            "10"
117          ]
118        },
119        {
120          "type": "categorical",
121          "categories": [
122            "-1",
123            "0",
124            "6",
125            "8",
126            "16",
127            "18"
128          ]
129        },
130        {
131          "type": "categorical",
132          "categories": [
133            "5",
134            "6",
135            "7",
136            "8",
137            "9"
138          ]
139        },
140        {
141          "type": "int",
142          "lower_bound": 0,
143          "upper_bound": 5
144        },
145        {
146          "type": "categorical",
147          "categories": [
148            "0",
149            "3",
```

```
150              "4",
151              "5",
152              "6",
153              "8",
154              "12"
155          ]
156        },
157        {
158          "type": "int",
159          "lower_bound": 0,
160          "upper_bound": 5
161        }
162      ],
163      "OS": [
164        {
165          "type": "real",
166          "lower_bound": -Infinity,
167          "upper_bound": Infinity
168        }
169      ]
170    },
171    "modeler": "Model_LCM",
172    "objective_id": 0,
173    "time": {
174      "tm_year": 2021,
175      "tm_mon": 1,
176      "tm_mday": 27,
177      "tm_hour": 21,
178      "tm_min": 52,
179      "tm_sec": 47,
180      "tm_wday": 2,
181      "tm_yday": 27,
182      "tm_isdst": 0
183    },
184    "uid": "0c5acce2-612d-11eb-8bf3-bbda784b918d"
185 }
```

Listing 2: Example Surrogate Model

4 User Interface

This section describes the essential APIs for a GPTune driver. For illustration purposes, we will describe the interfaces in the context of tuning the runtime of the parallel QR factorization routine $PDGEQRF$ from the ScaLAPACK package [5]. For the QR factorization time of a matrix, we can consider the matrix dimensions (m, n) to be the task parameters. Let row block size, column block size, MPI count per node , number of row processes, denoted by $(mb, nb, npernode, p)$ be the parameters to be tuned assuming fixed numbers of compute nodes and cores per node. Note that other arguments affecting the runtime such as the number of column processes and threads per process can be derived from these arguments. Finally let (r) be the QR runtime with fixed task parameters and tuning parameters.

4.1 Modify the application code

If the application code is not distributed-memory parallel, no modification is required; otherwise, GPTune replies on the MPI spawning approach to launch the application code, which requires slight modifications of the application code. For example assuming the application code is written in C language, the user needs to insert MPI_Comm_get_parent and MPI_Comm_disconnect after MPI_Init and before MPI_Finalize. The following example assumes that the parameters are passed in via the command line options (i.e., argv) and the function values are dumped into a file. For other options of passing parameters and returning objectives, see Section 3.2.2.

```
1 int main(int argc, char *argv[]){
2 MPI_Init( &argc, &argv );
3 MPI_Comm parent; MPI_Comm_get_parent(&parent); /* Get the inter communicator. */
4 .../* Read parameters from argv, compute the objective and dump it to a file. */
5 MPI_Comm_disconnect(&parent); /* Disconnect the inter communicator. */
6 MPI_Finalize();}
```

4.2 Define the objective function to be minimized

```
1 # define the multi or single objective function
2 def objectives(point):
3   # extract task and tuning parameters from "point", call the application code,
      and collect the results in "res".
4   return res
```

- **point** [Dict] (input): A dictionary containing the task parameter and tuning parameter arguments of the objective function. point[“I_j”] is the value of the task parameter argument named “I_j”, point[“P_j”] is the value of the parameter argument named “P_j”, see Section 4.7 for their definitions. The dictionary can also contain point[“C”] where “C” is a global constant (e.g. total node or core counts) passed to GPTune, see Section 4.8 for its definition.

- **res** [numpy array, shape=(1,γ)] (output): Array containing the objective function value, γ denotes dimension of the output space $\mathbb{OS}$.

QR EXAMPLE:
point[“m”], point[“n”], point[“mb”], point[“nb”], point[“npernode”], point[“p”] are the task parameter and tuning parameter arguments for $PDGEQRF$. The user is responsible for writing a driver code that dumps the parameters to an input file, calls $PDGEQRF$ and reads the runtime from an output file and returns the runtime in **res** with $\gamma = 1$, e.g., using a MPI spawning approach. See Section 5 for more details.

4.3 Define the performance models

```
1 # define the coarse performance models
2 def models(point):
3 # extract task and tuning parameters from "point", call the performance models,
      and collect the results in "res".
4 return res
```

- **point** [Dict] (input): A dictionary containing the task parameter and tuning parameter arguments of the objective function. point[“I_j”] is the value of the task parameter argument named “I_j”, point[“P_j”] is the value of the parameter argument named “P_j”, see Section 4.7 for their definitions. Note that the input “point” is exactly the same as that for “objectives”.

- **res** [numpy array, shape=$(1,\tilde{\gamma})$] (output): Array containing the performance models outputs. Note that the number of performance models $\tilde{\gamma}$ can be different from γ .

QR EXAMPLE:
point[“m”], point[“n”], point[“mb”], point[“nb”], point[“npernode”], point[“p”] are the task parameter and tuning parameter arguments for $PDGEQRF$. The user can use a simple performance model **res**$= [mn^2/npernode/nodes/PF]$, with PF denotes the peak flop rate of the machine.

4.4 Define the performance models update

```
1 # define the coarse performance models update function
2 def models_update ( data ):
3    for i in range ( len ( data.I )):
4    # update the hyperparameters data.D[i] of the performance model for each task
       uing data.I[i], data.O[i], data.P[i]
```

- **data** [Class] (input/output): The data class containing all tuning parameter configurations, function evaluations and possible hyperparameters for each task. See Section 4.11 for more details.

QR EXAMPLE:
For the performance model **res**$= [mn^2/npernode/nodes/PF]$, one can either use a constant PF, or a dynamic hyperparameter data.D[i][“PF”] per task i, which can be udpated using the growing sized samples.

4.5 Edit the meta JSON file (from command line)

For each application, the GPTune python driver requires a JSON file, located at ./.gptune/meta.json, defining application name, compute resources needed, and software dependence for both the current tuning experiment and loadable historical data. The meta.json file can be manually edited. Alternatively, the meta file can be automatically generated using jq from command line. For example:

```
1 tp=application_name      # define application name
2 nodes=1                        # define number of compute nodes
3 cores=4                        # define number of cores per node
4 machine=mymachine      # define machine name
5 proc=intel-i7                    # define processor type
6
7 # generate current and loadable software information. Change the software names
      and versions as needed.
8 software_json=$(echo ",\"software_configuration\":{\"openmpi\":{\"version_split\":
      [4,0,1]},\"scalapack\":{\"version_split\": [2,1,0]},\"gcc\":{\"version_split\":
      [8,3,0]}}")
9 loadable_software_json=$(echo ",\"loadable_software_configurations\":{\"openmpi\"
      :{\"version_split\": [4,0,1]},\"scalapack\":{\"version_split\": [2,1,0]},\"gcc\
      ":{\"version_split\": [8,3,0]}}")
```

```
10  # generate current and loadable machine information
11  machine_json=$(echo ",\"machine_configuration\":{\"machine_name\":\"$machine\",\"$
        proc\":{\"nodes\":$nodes,\"cores\":$cores}}")
12  loadable_machine_json=$(echo ",\"loadable_machine_configurations\":{\"$machine\":{\
        "$proc\":{\"nodes\":$nodes,\"cores\":$cores}}}")
13  # generate the meta file with jq
14  app_json=$(echo "{\"tuning_problem_name\":\"$tp\"")
15  echo "$app_json$machine_json$software_json$loadable_machine_json$
        loadable_software_json}" | jq '.' > .gptune/meta.json
```

4.6 Read the meta JSON file

```
1  (machine, processor, nodes, cores) = GetMachineConfiguration()
2  # read meta information defined in the ./gptune/meta.json file
```

4.7 Define the tuning parameter, task parameter and output spaces

```
1  IS = Space([I_1, I_2, ..., I_α]): # task parameter space with instances of supported spaces
        , α is the dimension of the task parameter space
2  PS = Space([P_1, P_2, ..., P_β]): # tuning parameter space with instances of supported
        spaces , β is the dimension of the tuning parameter space
3  OS = Space([O_1, O_2, ..., O_γ]): # output space with instances of supported spaces , γ is
        the dimension of the output space. Note that the return value of the Callable "
        objectives" is a point in this space.
4  constraints = {"cst1" : cst1, ...} # constraints for I_j and P_j
```

- I_j, P_j [Scikit-optimize spaces] (input): Please refer to `https://scikit-optimize.github.io/stable/modules/classes.html?highlight=space#module-skopt.space.space` for the scikit-optimize spaces. Supported spaces for I_j (and P_j) :
  I_j = Integer (low , high, transform="normalize", name="I_j"), here the "normalize" transform converts the integers in the range [low,high] to real numbers in $[0, 1]$, and vice versa. Please refer to `https://scikit-optimize.github.io/stable/modules/generated/skopt.space.transformers.Normalize.html` for more details about the transform. Note that it is required that low<high.
  I_j = Real(low, high, transform="normalize", name="I_j"), here the "normalize" transform converts real numbers in the range [low,high] to $[0, 1]$. It is required that low<high.
  I_j = Categoricalnorm (categories, transform="onehot", name="I_j"). Note: Categoricalnorm is compatible with Categorical but with a modified transform "onehot" that converts the categorical data to real numbers in $[0, 1)$. Spefically, the transform first converts the categorical data to its onehot encoding which is a $1 \times \alpha$ binary array with only element of 1 (e.g. the kth element), then converts it to the real number $(k - 1)/\alpha + 10^{-12}$. Conversely, the modified transform converts any number in $[(k - 1)/\alpha, k/\alpha)$ to the onehot encoding with kth element being 1, representing the kth category.

- O_j [Scikit-optimize spaces] (input): Supported spaces:
  O_j = Real(low, high, name="O_j"). Note: if low and high are unknown, they can be set to float("-Inf") and float("Inf"), respectively.

28

- **cst1** [string or function] (input): define one constraint using (a) a string, e.g. as cst1 = "$I_1 + P_2 + C < 10$", or (b) a callable function, e.g., as

```
1 # define a constraint function
2 def cst1(I_1, P_1, C):
3    return I_1 + P_2 + C < 10
```

Here I_1 and P_2 are task and tuning parameters, and C is a global constant (see Section 4.11).

QR EXAMPLE:

```
1 nodes=1
2 cores=4
3 bunit=8 # the block size is integer multiple of bunit
4 m   = Integer (128 , 2000, transform="normalize", name="m") # row dimension
5 n   = Integer (128 , 2000, transform="normalize", name="n") # column dimension
6 IS = Space([m, n]) # task parameter space
7 mb  = Integer (1 , 16, transform="normalize", name="mb") # row block size (divided
       by bunit)
8 nb  = Integer (1 , 16, transform="normalize", name="nb") # column block size (
       divided by bunit)
9 npernode  = Integer (0 , int(log2(cores)), transform="normalize", name="npernode")
       # (log2) of number of MPIs per node
10 p   = Integer (1 , nodes*cores, transform="normalize", name="p") # number of row
       processes
11 PS = Space([mb, nb, npernode, p]) # tuning parameter space
12 r   = Real (float("-Inf") , float("Inf"), name="r") # runtime
13 OS = Space([r]) # output space
14 cst1 = "mb*bunit*p<=m"
15 cst2 = "nb*bunit*nodes*2**npernode<=n*p"
16 cst3 = "nodes*2**npernode>=p"
17 constraints = {"cst1" : cst1, "cst2" : cst2, "cst3" : cst3} # constraints for task
       parameters and tuning parameters
18 constants={"nodes":nodes,"cores":cores,"bunit":bunit} # global constants
```

4.8   Define the tuning problem

```
1 problem = TuningProblem(IS, PS, OS, objectives, constraints, models, constants)
2 # define the tuning problem from the spaces, objective function, constraints,
       coarse performance models and constants
```

- **IS** [Space] (input): The task parameter space defining the objective function

- **PS** [Space] (input): The tuning parameter space defining the tuning parameters of the objective function

- **OS** [Space] (input): The output space defining the output of the objective function

- **objectives** [Callable] (input): The objective function to be minimized

- **constraints** [Dict] (input): The constraints for the task parameters and tuning parameters

- **models** [Callable] (input): The avaiable performance models. This can be None if no model is available.

- **constants** [Dict] (input): The global constants that can be used in objectives, constraints and models.

- **problem** [Class] (output): The tuning problem

4.9 Define the computation resource

```
1 computer = Computer(nodes, cores, hosts) # define the computation resource used in
      GPTune's internal computation. Note: the same amount of resource can be used
     for invoking the objective functions as the function evaluation and the tuner's
      internal computation do not run parallel to each other. See Section 5 for
     examples.
```

- **nodes** [Int] (input): The minimum number of MPI processes used for GPTune different phases

- **cores** [Int] (input): The maximum number of threads per MPI process used for GPTune different phases

- **hosts** [Collection] (input): The list of hostnames

- **computer** [Class] (output): The computer class

4.10 Define and validate the options

```
1 options = Options() # define the default options
```

- **options** [Class] (output): The options class. See GPTune/options.py and Section 4.19 for all the options.

```
1 options['item'] = val # set the option named 'item' to val.
```

```
1 options.validate(computer) # check the options specified by the user
```

- **computer** [Class] (input): The computer class

4.11 Create the data class for storing samples of the spaces

```
1 data = Data(problem,I,P,O,D) # define the data class
```

- **problem** [Class] (input): The tuning problem class

- **data** [Class] (output): The data class to be used for storing sampled tasks, tuning parameters and outputs.

  data.I=I [list of lists]: Each entry is a list of length α representing one task sample.

  data.P=P [list of [list of lists]]: Each entry corresponds to one task sample. For each entry (list of lists), each entry is a list of length β representing one tuning parameter configuration.

  data.O=O [list of numpy arrays]: Each entry is a numpy array of shape $(,\gamma)$ representing objective function values for one task sample. One row of each array represents the objective function values for one tuning parameter configuration.

  data.D=D [list of Dicts]: Each entry is a dictionary containing constants/hyperparameters for data of one task sample.

  On return, data.I=None, data.P=None, data.O=None, data.D=None if I,P,O,D is not provided. See Section 4.13 for details.

4.12 Initialize GPTune

```
gptune = GPTune(problem, computer, data, options, driverabspath, models_update)
# initialize the tuner from the meda data
```

- **problem** [Class] (input): The tuning problem

- **computer** [Class] (input): The computer

- **data** [Class] (input): The data class. If any of data.I, data.P and data.O is None, the tuner will generate random samples for it later.

- **options** [Class] (input): The options class.

- **driverabspath** [string] (input): Absolute path of the file containing this call. Default to None. We recommend passing driverabspath=os.path.abspath(__file__).

- **models_update** [Callable] (input): The coarse performance update function (see Section 4.4). This can be set to None if there is no performance model, or the performance model requires no update.

- **gptune** [Class] (input): The tuner class that registers problem, computer, data and options as gptune.problem, gptune.computer, gptune.data, and gptune.options, respectively.

4.13 Call multi-task learning algorithm (MLA)

```
(data, models,stats) = gptune.MLA(NS, NI, Igiven, NS1)
# build the MLA models and search for the optimum tuning parameters on each task
```

- **gptune** [Class]: The tuner

- **NS** [Int] (input): Number of total samples per task to be returned. Note that the tuner returns immediately if the number of samples in historical data is more than NS.

- **NI** [Int] (input): Number of tasks to be modeled (i.e., NI=δ) (Note that in the MLA interface NI should match the number of tasks in the historical data (if present). If one needs to add new tasks to the historical data, use the TLA interface in Section 4.14 instead.)

- **Igiven** [list of lists] (input): A list of prescribed task parameters. Note that Igiven should match the list of task parameters in the historical data (if present).

- **NS1** [Int] (input):If no historical data is presesnt, the tuner generates NS1 initial random tuning parameter samples before starting the adaptive model refinement.

- **data** [Class] (output): The data class containing all the task, tuning parameter and output sampled by the tuner. len(data.I)=NI, len(data.P)=NI, len(data.O)=NI, and len(data.P[0])=NS.

- **models** [list of Class] (output): Each entry represents the trained LCM model for one objective in the adaptive model refinement.

- **stats** [Dict] (output): Memory and CPU profiles generated by the tuner

4.14   Call transfer learning algorithm (TLA)

```
1 (aprxopts,objval,stats) = gptune.TLA(newtask, NS)
2 # Use existing data and MLA model on pre-tuned tasks (stored in gptune.data and
      gptune.models) and TLA to search for the optimum tuning parameters on each new
      task
```

- **gptune** [Class]: The tuner that encapsulates the previous tuning data and models.

- **newtask** [list of lists] (input): newtask consists of list of prescribed tasks to be tuned

- **NS** [Int] (input): Maximum number of objective function evaluations

- **aprxopts** [list of lists] (output): Each entry (list) of the list corresponds to the predicted best tuning parameters

- **objval** [list of numpy arrays] (output): Each entry of the list corresponds to objective function values using the predicted tuning parameters for one task

- **stats** [Dict] (output): Memory and CPU profiles generated by the tuner

4.15   Call opentuner

```
1 (data,stats)=OpenTuner(T, NS, tp, computer, run_id)
2 # initialize and call opentuner to generate objective function samples. If there
      are multiple tasks, opentuner is invoked one task each time.
```

- **T** [list of lists] (input): A list of prescribed task parameters of length NI.

- **NS** [Int] (input): Number of total samples per task to be returned.

- **tp** [Class] (output): The tuning problem defined in Section 4.8.

- **computer** [Class] (input): The computer.

- **runid** [String] (input): Name of the tuner (default to "OpenTuner").

- **data** [Class] (output): The data class containing all the task, tuning parameter and output sampled by the tuner. len(data.I)=NI, len(data.P)=NI, len(data.O)=NI, and len(data.P[0])=NS.

- **stats** [Dict] (output): Memory and CPU profiles generated by the tuner

4.16 Call hpbandster

```
1 (data,stats)=HpBandSter(T, NS, tp, computer, run_id)
2 # initialize and call opentuner to generate objective function samples. If there
    are multiple tasks, opentuner is invoked one task each time.
```

- **T** [list of lists] (input): A list of prescribed task parameters of length NI.

- **NS** [Int] (input): Number of total samples per task to be returned.

- **tp** [Class] (output): The tuning problem defined in Section 4.8.

- **computer** [Class] (input): The computer.

- **runid** [String] (input): Name of the tuner (default to "HpBandSter").

- **data** [Class] (output): The data class containing all the task, tuning parameter and output sampled by the tuner. len(data.I)=NI, len(data.P)=NI, len(data.O)=NI, and len(data.P[0])=NS.

- **stats** [Dict] (output): Memory and CPU profiles generated by the tuner

4.17 Invoke GPTune (from command line): default mode

```
1 mpirun --oversubscribe --allow-run-as-root --mca pmix_server_max_wait 3600 --mca
    pmix_base_exchange_timeout 3600 --mca orte_abort_timeout 3600 --mca
    plm_rsh_no_tree_spawn true -n 1  python ./application_tuner_driver.py
```

Here application_tuner_driver.py is the GPTune python driver that contains definitions from Section 4.2 to Section 4.16. The above command invokes the GPTune tuning process. Note that these MPI runtime parameters are necessary for OpenMPI 4.0.1, higher versions have not been tested extensively.

4.18 Invoke GPTune (from command line): RCI mode

```
1 python ./application_tuner_driver_rci.py
```

Here application_tuner_driver_rci.py is the GPTune python driver similar to application_tuner_driver.py, which contains definitions from Section 4.3 to Section 4.14. For RCI mode, the above command will execute phases in GPTune as usual, but without calling the application code. Instead, when function evaluation is needed, the required samples will be stored in the database ./gptune_db/application_name.json, which is located in the same directory as application_tuner_driver_rci.py. The user can then search for required samples in the database, invoke the application code in bash,

write the evaluation results into the database, and then call the GPTune python driver again to ask for next samples. As such, there is no need to define objective function as in Section 4.2 or modify the application code as in Section 4.1. Moreover, OpenMPI and its runtime parameters are not mandatory (e.g., MPICH or Spectrum MPI can also be used). Note that in application_tuner_driver_rci.py, options['RCI_mode']=True is required. In addition, the interface to opentuner (Section 4.15) or hpbandster (Section 4.16) is not supported in RCI mode. See Section 5.3 for a complete example.

4.19   GPTune options

The options affecting the efficiency of GPTune are listed below.

```python
class Options(dict):
  def __init__(self, **kwargs):

    """ Options for GPTune """
    RCI_mode = False           # whether the reverse communication mode will be used
    mpi_comm = None            # The mpi communiator that invokes gptune if mpi4py
    is installed
    distributed_memory_parallelism = False   # Using
    distributed_memory_parallelism for the modeling (one MPI per model restart) and
     search phase (one MPI per task)
    shared_memory_parallelism      = False   # Using shared_memory_parallelism for
     the modeling (one thread per model restart) and search phase (one thread per
    task)
    verbose = False      # Control the verbosity level
    oversubscribe = False      # Set this to True when the physical core count is
    less than computer.nodes*computer.cores and the --oversubscribe MPI runtime
    option is used

    """ Options for the function evaluation """
    objective_evaluation_parallelism   = False  # Using
    distributed_memory_parallelism or shared_memory_parallelism for evaluating
    multiple application instances in parallel
    objective_multisample_processes = None  # Number of MPIs each handling one
    application call
    objective_multisample_threads = None  # Number of threads each handling one
    application call
    objective_nprocmax = None # Maximum number of cores for each application call,
     default to computer.cores*computer.nodes-1

    """ Options for the sampling phase """
    sample_class = 'SampleLHSMDU' # Supported sample classes: 'SampleLHSMDU', '
    SampleOpenTURNS'
    sample_algo = 'LHS-MDU' # Supported sample algorithms: 'LHS-MDU' --Latin
    hypercube sampling with multidimensional uniformity, 'MCS' --Monte Carlo
    Sampling
    sample_max_iter = 10**9  # Maximum number of iterations for generating random
    samples and testing the constraints

    """ Options for the modeling phase """
    model_class = 'Model_LCM' # Supported sample algorithms: 'Model_GPy_LCM' --
    LCM from GPy, 'Model_LCM' -- LCM with fast and parallel inversion
    model_threads = None  # Number of threads used for building one GP model in
    Model_LCM
```

```
26    model_processes = None # Number of MPIs used for building one GP model in
      Model_LCM
27    model_restarts = 1 # Number of random starts each building one initial GP
      model
28    model_restart_processes = None  # Number of MPIs each handling one random
      start
29    model_restart_threads = None   # Number of threads each handling one random
      start
30    model_max_iters = 15000   # Number of maximum iterations for the optimizers
31    model_latent = None # Number of latent functions for building one LCM model,
      defaults to number of tasks
32
33    """ Options for the search phase """
34    search_threads = None  # Number of threads in each thread group handling one
      task
35    search_processes = 1  # Reserved option
36    search_multitask_threads = None # Number of threads groups each handling one
      task
37    search_multitask_processes = None # Number of MPIs each handling one task
38    search_algo = 'pso' # Supported search algorithm in pygmo: single-objective: '
      pso' -- particle swarm, 'cmaes' -- covariance matrix adaptation evolution.
      multi-objective 'nsga2' -- Non-dominated Sorting GA, 'nspso' -- Non-dominated
      Sorting PSO, 'maco' -- Multi-objective Hypervolume-based ACO, 'moead' -- Multi-
      objective EA vith Decomposition
39    search_pop_size = 1000 # Population size in pgymo
40    search_gen = 1000   # Number of evolution generations in pgymo
41    search_evolve = 10  # Number of times migration in pgymo
42    search_max_iters = 10   # Max number of searches to get results respecting the
      constraints
43    search_more_samples = 1   # Maximum number of points selected using a multi-
      objective search algorithm
```

Listing 3: Default GPTune options.

5 Example code

5.1 ScaLAPACK QR

As mentioned in Section 4, we can use (m, n) to define a task, then $(mb, nb, npernode, p)$ to define the tuning parameters. Here we use simplified codes to illustrate a few typical use cases with GPTune. Please refer to GPTune/examples/ Scalapack-DPDGEQRF/scalapack_MLA.py, GPTune/examples/Scalapack-DPDGEQRF/scalapack-driver/spt/pdqrdriver.py, and GPTune/examples/ Scalapack-DPDGEQRF/scalapack-driver/src/pdqrdriver.f for the complete working codes. Use GPTune/run_examples.sh to run the tests.

5.1.1 Preparing the meta JSON file

For each application directory, GPTune requires the use of a meta JSON file located at ./gptune/meta.json to define application names, machine, software configurations for the current experiment, as well as those for usable historical/shared dababases. For example, a meta JSON file for *PDGEQRF* looks like the following (Listing 4)

```json
{
  "tuning_problem_name": "PDGEQRF",
  "machine_configuration": {
    "machine_name": "mymachine",
    "intel": {
      "nodes": 1,
      "cores": 16
    }
  },
  "software_configuration": {
    "openmpi": {
      "version_split": [
        4,
        0,
        1
      ]
    },
    "scalapack": {
      "version_split": [
        2,
        1,
        0
      ]
    },
    "gcc": {
      "version_split": [
        8,
        3,
        0
      ]
    }
  },
  "loadable_machine_configurations": {
    "mymachine": {
      "intel": {
        "nodes": 1,
        "cores": 16
      }
    }
  },
  "loadable_software_configurations": {
    "openmpi": {
      "version_split": [
        4,
        0,
        1
      ]
    },
    "scalapack": {
      "version_split": [
        2,
        1,
        0
      ]
    },
    "gcc": {
```

```
57        "version_split": [
58            8,
59            3,
60            0
61        ]
62      }
63   }
64 }
```

<div align="center">Listing 4: meta.json for PDGEQRF</div>

This file can be manually edited for the current tuning experiment.

Alternatively, this process can be automated by using the command-line JSON parser jq. For example:

```
1 tp=PDGEQRF
2 nodes=1
3 cores=16
4 machine=mymachine
5 proc=intel
6 software_json=$(echo ",\"software_configuration\":{\"openmpi\":{\"version_split\":
      [4,0,1]},\"scalapack\":{\"version_split\": [2,1,0]},\"gcc\":{\"version_split\":
      [8,3,0]}}")
7 loadable_software_json=$(echo ",\"loadable_software_configurations\":{\"openmpi\"
      :{\"version_split\": [4,0,1]},\"scalapack\":{\"version_split\": [2,1,0]},\"gcc\
      ":{\"version_split\": [8,3,0]}}")
8 machine_json=$(echo ",\"machine_configuration\":{\"machine_name\":\"$machine\",\"$
      proc\":{\"nodes\":$nodes,\"cores\":$cores}}")
9 loadable_machine_json=$(echo ",\"loadable_machine_configurations\":{\"$machine\":{\
      "$proc\":{\"nodes\":$nodes,\"cores\":$cores}}}")
10
11 app_json=$(echo "{\"tuning_problem_name\":\"$tp\"")
12 echo "$app_json$machine_json$software_json$loadable_machine_json$
      loadable_software_json}" | jq '.' > .gptune/meta.json
```

5.1.2   MLA

This example builds a LCM model of the *PDGEQRF* example for two user specified tasks [[400,500],[800,600]]. The Python interface to the Fortran application (pdqrdriver) will dump the task parameter and tuning parameters into an input file named ''GPTune/examples/scalapack-driver/exp/MACHINE_NAME/TUNER_NAME/JOBID/QR.in'', invoke the Fortran application code pdqrdriver.f, then read the return values from an output file named ''GPTune/examples/scalapack-driver/exp/MACHINE_NAME/TUNER_NAME/JOBID/QR.out''. The variables ''MACHINE_NAME'', ''TUNER_NAME'', and ''JOBID'' can be defined by the user.

   In terms of computation resource, "nodes=1" and "cores=16" are used for GPTune's several phases and for invoking the application code (pdqrdriver.f). As we use MPI spawn to invoke the application, one process is reserved as the spawning process. Therefore it's recommended allocating nodes+1 compute nodes with 1 node reserved for MPI spawning.

   Note that to reduce the runtime noise, the Python driver will execute the same task and tuning parameter configuration three times (niter=3 as an argument of pdqrdriver) and return the minimum runtime as the function value.

```
1
```

```python
2  ''' Pass the inputs and parameters from Python to the Fortran driver using files
      RUNDIR/QR.in, note that the the inputs and parameters are duplicated for niter
      times.'''
3  def write_input(params, RUNDIR, niter=1):
4    fin = open("%s/QR.in"%(RUNDIR), 'w')
5    fin.write("%d\n"%(len(params) * niter))
6    for param in params:
7      for k in range(niter):
8        fin.write("%2s%6d%6d%6d%6d%6d%20.13E\n"%(param[0], param[1], param[2],
      param[5], param[6], param[9], param[10],param[11]))
9    fin.close()
10
11 ''' Execute the Fortran driver using MPI spawn. Note that the inputs and
      parameters are passed to Frotran using environment vairables, command lines and
        files. '''
12 def execute(nproc, nthreads, RUNDIR):
13   info = MPI.Info.Create()
14   info.Set('env', 'OMP_NUM_THREADS=%d\n' %(nthreads))
15   print('exec', "%s/pdqrdriver"%(BINDIR), 'args', "%s/"%(RUNDIR), 'nproc', nproc)
16   comm = MPI.COMM_SELF.Spawn("%s/pdqrdriver"%(BINDIR), args="%s/"%(RUNDIR),
      maxprocs=nproc,info=info)
17   comm.Disconnect()
18   return
19
20 ''' Read the runtime from the output file RUNDIR/QR.out which contains the runtime
       for the same parameters by running QR factorization for niter times. Only the
      minimum among the niter runtimes is returned. '''
21 def read_output(params, RUNDIR, niter=1):
22   fout = open("%s/QR.out"%(RUNDIR), 'r')
23   times = float('Inf')
24   for line in fout.readlines():
25     words = line.split()
26     if (len(words) > 0 and words[0] == "WALL"):
27       if (words[9] == "PASSED"):
28         mytime = float(words[7])
29         if (mytime < times):
30           times = mytime
31   fout.close()
32   return times
33
34 ''' The Python dirver that writes parameters to individual files, runs the Fortran
       driver, and read the runtime from individual files. Note the same parameter is
       executed niter times. '''
35 def pdqrdriver(params, niter=10,JOBID: int = None):
36   global EXPDIR  # path to the input and output files
37   global BINDIR  # path to the executable
38   global ROOTDIR # path to the folder "scalapack-driver"
39
40   ROOTDIR = os.path.abspath(os.path.join(os.path.realpath(__file__), '/scalapack-
      driver'))
41   BINDIR = os.path.abspath(os.path.join(ROOTDIR, "bin", MACHINE_NAME))
42   EXPDIR = os.path.abspath(os.path.join(ROOTDIR, "exp", MACHINE_NAME + '/' +
      TUNER_NAME))
43
44   if (JOBID==-1):  # -1 is the default value if jobid is not set
45     JOBID = os.getpid()
```

```
46    RUNDIR = os.path.abspath(os.path.join(EXPDIR, str(JOBID)))
47    os.system("mkdir -p %s"%(EXPDIR))
48    os.system("mkdir -p %s"%(RUNDIR))
49    idxproc = 8  # the index in params representing MPI counts
50    idxth = 7    # the index in params representing thread counts
51    write_input(params, RUNDIR, niter=niter)
52    execute(params[idxproc], params[idxth], RUNDIR)
53    times = read_output(params, RUNDIR, niter=niter)
54    return times
55
56  ''' The objective function required by GPTune. '''
57  def objectives(point):
58    # global constants defined in TuningProblem
59    nodes = point['nodes']
60    cores = point['cores']
61    bunit = point['bunit']
62    # task and tuning parameters
63    m = point['m']
64    n = point['n']
65    mb = point['mb']*bunit
66    nb = point['nb']*bunit
67    p = point['p']
68    npernode = 2**point['npernode']
69    nproc = nodes*npernode
70    nthreads = int(cores / npernode)
71    if(nproc==0 or p==0 or nproc<p): # this become useful when the parameters
        returned by TLA1 do not respect the constraints
72      print('Warning: wrong parameters for objective function!!!')
73      return 1e12
74    q = int(nproc / p)
75    nproc = p*q
76    params = [('QR', m, n, nodes, cores, mb, nb, nthreads, nproc, p, q, 1.)]
77    elapsedtime = pdqrdriver(params, niter = 3)
78    print(params, ' scalapack time: ', elapsedtime)
79    return elapsedtime
80
81  def main():
82    global JOBID
83    mmax = 2000  # maximum row dimension
84    nmax = 2000  # maximum column dimension
85    ntask = 2   # 2 tasks used for MLA
86    NS = 20   # 20 samples per task
87    JOBID = 0    # JOBID is part of the input/output file names
88    TUNER_NAME='GPTune' # TUNER_NAME is part of the input/output file names
89    # read the meta.json file
90    (machine, processor, nodes, cores) = GetMachineConfiguration()
91    print ("machine: " + machine + " processor: " + processor + " num_nodes: " + str
        (nodes) + " num_cores: " + str(cores))
92    os.environ['MACHINE_NAME'] = machine
93    os.environ['TUNER_NAME'] = TUNER_NAME
94
95    """ Define and print the spaces and constraints """
96    # Task Parameters
97    m     = Integer (128 , mmax, transform="normalize", name="m")
98    n     = Integer (128 , nmax, transform="normalize", name="n")
99    IS = Space([m, n])
```

```python
100    # Tuning Parameters
101    mb     = Integer (1 , 512, transform="normalize", name="mb")
102    nb     = Integer (1 , 512, transform="normalize", name="nb")
103    npernode      = Integer     (0, int(math.log2(cores)), transform="normalize",
         name="npernode")
104    p      = Integer (1 , nodes*cores, transform="normalize", name="p")
105    PS = Space([mb, nb, npernode, p])
106    # Output
107    r      = Real     (float("-Inf") , float("Inf"), name="r")
108    OS = Space([r])
109    # Constraints
110    cst1 = "mb*bunit*p<=m"
111    cst2 = "nb*bunit*nodes*2**npernode<=n*p"
112    cst3 = "nodes*2**npernode>=p"
113    constraints = {"cst1" : cst1, "cst2" : cst2, "cst3" : cst3}
114    constants={"nodes":nodes,"cores":cores,"bunit":bunit}
115    print(IS, PS, OS, constraints)
116
117    problem = TuningProblem(IS, PS, OS, objective, constraints, None, constants)
118    computer = Computer(nodes = nodes, cores = cores, hosts = None)
119
120    """ Set and validate options """
121    options = Options()
122    options['model_restarts'] = 4   # number of GP models being built in one
         iteration (only the best model is retained)
123    options['distributed_memory_parallelism'] = True # True: Use MPI. One MPI per
         model start in the modeling phase, one MPI per task in the search phase
124    options['shared_memory_parallelism'] = False # True: Use threads. One thread per
          model start in the modeling phase, one MPI per task in the search phase
125    options.validate(computer = computer)
126
127    """ Intialize the tuner with existing data"""
128    data = Data(problem)  # intialize with empty data, but can also load data from
         previous runs
129    gptune = GPTune(problem, computer = computer, data = data, options = options)
130
131    """ Building MLA with the given list of tasks """
132    giventask = [[400,500],[800,600]]
133    NI = len(giventask)
134    (data, models,stats) = gptune.MLA(NS=NS, NI=NI, Igiven =giventask, NS1 = max(NS
         //2,1))
135    print("stats: ",stats)
136    """ Print all input and parameter samples """
137    for tid in range(NI):
138      print("tid: %d" % (tid))
139      print("    m:%d n:%d" % (data.I[tid][0], data.I[tid][1]))
140      print("    Ps ", data.P[tid])
141      print("    Os ", data.O[tid].tolist())
142      print('    Popt ', data.P[tid][np.argmin(data.O[tid])], 'Oopt ', min(data.O[
         tid])[0], 'nth ', np.argmin(data.O[tid]))
143
144 if __name__ == "__main__":
145    main()
```

Listing 5: scalapack QR example: scalapack_MLA.py.

```fortran
      PROGRAM PDQRDRIVER
      CHARACTER*200        FILEDIR
      INTEGER              IAM
      INTEGER              master
      PARAMETER            ( NIN = 1, NOUT = 2)
      DATA                 KTESTS, KPASS, KFAIL, KSKIP /4*0/

*     Get starting information
      CALL GETARG(1,FILEDIR)
      CALL MPI_INIT(ierr)
      CALL MPI_COMM_GET_PARENT(master, ierr)
      CALL BLACS_PINFO( IAM, NPROCS )

*     Open input file
      OPEN( NIN, FILE=trim(FILEDIR)//'QR.in', STATUS='OLD' )
      IF( IAM.EQ.0 ) THEN
          OPEN( NOUT, FILE=trim(FILEDIR)//'QR.out', STATUS='UNKNOWN' )
      END IF

*     Read number of configurations
      READ( NIN, FMT = 1111 ) NBCONF
*     Print headings
      IF( IAM.EQ.0 ) THEN
          WRITE( NOUT, FMT = * )
          WRITE( NOUT, FMT = 9995 )
          WRITE( NOUT, FMT = 9994 )
          WRITE( NOUT, FMT = * )
      END IF

*     Run the computation for NBCONF times and dump the runtime to QR.out
      DO 50 CONFIG = 1, NBCONF
*        ...
*        ...
      END DO

*     Print out ending messages and close output file
      IF( IAM.EQ.0 ) THEN
         KTESTS = KPASS + KFAIL + KSKIP
         WRITE( NOUT, FMT = * )
         WRITE( NOUT, FMT = 9992 ) KTESTS
         IF( CHECK ) THEN
             WRITE( NOUT, FMT = 9991 ) KPASS
             WRITE( NOUT, FMT = 9989 ) KFAIL
         ELSE
             WRITE( NOUT, FMT = 9990 ) KPASS
         END IF
         WRITE( NOUT, FMT = 9988 ) KSKIP
         WRITE( NOUT, FMT = * )
         WRITE( NOUT, FMT = * )
         WRITE( NOUT, FMT = 9987 )
      END IF

*     Close input and output files
      CLOSE( NIN )
      IF( IAM.EQ.0 ) THEN
          CLOSE( NOUT )
```

```
57        END IF
58
59 *    Disconnect the inter communicator, destroy BLACS grid and the inter
        communicator
60       call MPI_COMM_DISCONNECT(master, ierr)
61       CALL BLACS_EXIT( 1 )
62       call MPI_Finalize(ierr)
63
64 *    Formats
65  1111 FORMAT( I6 )
66  9995 FORMAT( 'TIME       M       N  MB  NB     P     Q Fact Time ','    MFLOPS
        CHECK  Residual' )
67  9994 FORMAT( '---- ------ ------ --- --- ----- ----- --------- ','-----------
        ------  --------' )
68  9992 FORMAT( 'Finished ', I6, ' tests, with the following results:' )
69  9991 FORMAT( I5, ' tests completed and passed residual checks.' )
70  9990 FORMAT( I5, ' tests completed without checking.' )
71  9989 FORMAT( I5, ' tests completed and failed residual checks.' )
72  9988 FORMAT( I5, ' tests skipped because of illegal input values.' )
73  9987 FORMAT( 'END OF TESTS.' )
74
75       STOP
76       END
```

Listing 6: pdqrdriver.f.

In order to run the tuning experiment, use the following mpirun command:

```
1 mpirun --oversubscribe --allow-run-as-root --mca pmix_server_max_wait 3600 --mca
    pmix_base_exchange_timeout 3600 --mca orte_abort_timeout 3600 --mca
    plm_rsh_no_tree_spawn true -n 1  python ./scalapack_MLA.py
```

Note that these MPI runtime parameters are necessary for OpenMPI 4.0.1, higher versions have
not been tested extensively.

As all the function evaluation data are check-pointed using JSON files (./gptune_db/applicationname.json),
GPTune is fault resilient. In case the tuning is interrupted due to timeout or code crashes, just
rerun the above command and GPTune will continue tuning as expected.

The following example runlog illustrates how to understand the code generated information.
First, the code prints out the IS, PS and OS. Then it shows the parallelization parameters being
used by GPTune. Next, the code prints different phases in each MLA iteration. Next, the tuner
runtime profile is printed. Finally, all and the best samples of each task are listed.

```
1
2 IS: Space([Integer(low=128, high=2000),
3     Integer(low=128, high=2000)])
4 PS: Space([Integer(low=1, high=128),
5     Integer(low=1, high=128),
6     Integer(low=4, high=31),
7     Integer(low=1, high=31)])
8 OS: Space([Real(low=-inf, high=inf, prior='uniform', transform='identity')])
9 constraints: {'cst1': 'mb * p <= m', 'cst2': 'nb * nproc <= n * p', 'cst3': 'nproc
    >= p'}
10
11
12 ------Validating the options
13   total core counts provided to GPTune: 32
```

42

```
14     ---> distributed_memory_parallelism: True
15     ---> shared_memory_parallelism: False
16     ---> objective_evaluation_parallelism: False
17
18   total core counts for modeling: 29
19    ---> model_processes: 6
20    ---> model_threads: 1
21    ---> model_restart_processes: 4
22    ---> model_restart_threads: 1
23
24   total core counts for search: 32
25    ---> search_processes: 1
26    ---> search_threads: 1
27    ---> search_multitask_processes: 31
28    ---> search_multitask_threads: 1
29
30   total core counts for objective function evaluation: 32
31    ---> core counts in a single application run: 31
32    ---> objective_multisample_processes: 1
33    ---> objective_multisample_threads: 1
34
35
36 ------Starting MLA with 2 tasks
37 MLA initial sampling:
38 ...
39 MLA iteration:  0
40 ...
41 MLA iteration:  9
42 ...
43 stats:  {'time_total': 61.6, 'time_fun': 18.4, 'time_search': 11.9, 'time_model':
    31.1}
44
45 tid: 0
46     m:400 n:500
47
48     Ps  [[28, 47, 15, 9], [9, 124, 26, 17], [38, 50, 15, 10], [1, 117, 28, 18],
49     [77, 73, 9, 3],[31, 50, 16, 7], [8, 117, 28, 19], [71, 67, 8, 2],
50     [35, 39, 15, 9], [13, 116, 27, 16],[121, 102, 7, 2], [101, 128, 5, 2],
51     [21, 13, 27, 3], [15, 127, 6, 2], [16, 14, 30, 2],[49, 30, 27, 2],
52     [106, 6, 29, 2], [108, 31, 6, 2], [32, 14, 31, 2], [86, 5, 25, 2]]
53
54     Os  [[0.0092][0.0157][0.01][0.0152][0.0083][0.0082][0.0171][0.0079]
55         [0.0104][0.0143][0.009][0.0112][0.0068][0.0102][0.006][0.0057]
56         [0.0066][0.0061][0.0059][0.0062]]
57
58     Popt  [49, 30, 27, 2] Oopt  0.005727
59
60 tid: 1
61     m:800 n:600
62
63     Ps  [[32, 40, 16, 8], [13, 116, 28, 17], [76, 65, 7, 2], [116, 5, 30, 6],
64     [33, 51, 15, 9], [5, 123, 26, 16], [65, 69, 9, 3], [116, 11, 29, 4],
65     [31, 44, 17, 7], [11, 117, 28, 18],[119, 124, 26, 6], [12, 72, 21, 8],
66     [87, 10, 31, 4], [1, 101, 7, 4], [122, 7, 22, 4],[40, 26, 28, 2],
67     [101, 21, 26, 2], [97, 37, 12, 3], [2, 2, 29, 2], [29, 10, 30, 1]]
68
```

```
69    Os    [[0.0136][0.0239][0.0174][0.0177][0.0141][0.0231][0.0174][0.014]
70          [0.0116][0.0248][0.0237][0.015 ][0.0144][0.0169][0.0145][0.0122]
71          [0.0118][0.0127][0.0201][0.0103]]
72
73    Popt  [29, 10, 30, 1] Oopt  0.01032
```

Listing 7: runlog of MLA

5.1.3   MLA+TLA

Inserting the following code segments after line 142 of Listing 5, GPTune will call TLA to predict the optimal tuning parameters for 1 new task [[450,450]].

```
1 """ Call TLA for 1 new task using the constructed LCM model"""
2 newtask = [[450, 450]]
3 (aprxopts, objval, stats) = gt.TLA1(newtask, NS=None)
4 print("stats: ", stats)
5
6 """ Print the optimal parameters and function evaluations"""
7 for tid in range(len(newtask)):
8   print("new task: %s" % (newtask[tid]))
9   print('    predicted Popt: ', aprxopts[tid], ' objval: ', objval[tid])
```

Listing 8: scalapack QR example: MLA+TLA.

The following runlog illustrates how the output of TLA looks like. Please refer to Section 5.1.2 for the output of MLA.

```
1
2 IS: Space([Integer(low=128, high=2000),
3     Integer(low=128, high=2000)])
4 PS: Space([Integer(low=1, high=128),
5     Integer(low=1, high=128),
6     Integer(low=4, high=31),
7     Integer(low=1, high=31)])
8 OS: Space([Real(low=-inf, high=inf, prior='uniform', transform='identity')])
9 constraints: {'cst1': 'mb * p <= m', 'cst2': 'nb * nproc <= n * p', 'cst3': 'nproc
      >= p'}
10
11 ------Validating the options
12   total core counts provided to GPTune: 32
13    ---> distributed_memory_parallelism: True
14    ---> shared_memory_parallelism: False
15    ---> objective_evaluation_parallelism: False
16
17   total core counts for modeling: 3
18    ---> model_processes: 1
19    ---> model_threads: 1
20    ---> model_restart_processes: 1
21    ---> model_restart_threads: 1
22
23   total core counts for search: 32
24    ---> search_processes: 1
25    ---> search_threads: 1
26    ---> search_multitask_processes: 31
27    ---> search_multitask_threads: 1
```

```
28
29    total core counts for objective function evaluation: 32
30     ---> core counts in a single application run: 31
31     ---> objective_multisample_processes: 1
32     ---> objective_multisample_threads: 1
33
34
35  ------Starting MLA with 10 tasks
36  ...
37
38  ------Starting TLA1 for task:  [[400, 500], [800, 600]]
39  ...
40  stats:  {'time_total': 1.337391381, 'time_fun': 0.906422981}
41  new task: [400, 500]
42      predicted Popt:  [38, 15, 27, 1]  objval:  [[0.004]]
43  new task: [800, 600]
44      predicted Popt:  [35, 15, 30, 1]  objval:  [[0.0103]]
```

Listing 9: runlog of MLA+TLA (the part of MLA is skipped)

5.2 SuperLU_DIST

For a typical SuperLU_DIST driver, one can use a given sparse matrix to define a task, and consider (COLPERM, LOOKAHEAD, npernode, nprows, NSUP, NREL) to be the tuning parameters affecting the objective function (e.g., runtime). Here COLPERM is the column permutation option, LOOKAHEAD is the size of the lookahead window in the factorization, npernode is the MPI count per compute node, nprows is the number of row processes, NSUP is the maximum supernode size, and NREL is the supernode relaxation parameter. Just like the ScaLAPACK example, we use a MPI spawn approach to invoke the application code (pddrive_spawn.c). However, this example does not use files for passing information between the driver and the application. Instead, some task parameters and tuning parameters are passed to the application through command lines, while the others are passed through environment variables; the output is passed back using the spawned MPI communicator.

5.2.1 Preparing the meta JSON file

Just like *PDGEQRF*, a meta JSON file located at ./gptune/meta.json needs to be edited for SuperLU_DIST.

5.2.2 MLA+TLA

The following example first calls MLA to build a LCM model for the SuperLU_DIST driver pddrive_spawn.c using two tasks [["g4.rua"], ["g20.rua"]], then calls TLA to predict the optimal tuning parameters for a new task [["big.rua"]]. Note that only the simplified code is shown here, please refer to GPTune/examples/SuperLU_DIST/superlu_MLA.py and GPTune/examples/SuperLU_DIST/superlu_dist/EXAMPLE/pddrive_spawn.c for the complete working codes.

```
1  def objectives(point):
2    nodes = point['nodes']
3    cores = point['cores']
```

45

```
4   RUNDIR = os.path.abspath(__file__ + "/../superlu_dist/build/EXAMPLE") # the path
       to the executable
5   INPUTDIR = os.path.abspath(__file__ + "/../superlu_dist/EXAMPLE/") # the path to
       the matrix collection
6
7   matrix = point['matrix']
8   COLPERM = point['COLPERM']
9   LOOKAHEAD = point['LOOKAHEAD']
10  nprows = point['nprows']
11  npernode = 2**point['npernode']
12  nproc = nodes*npernode
13  nthreads = int(cores / npernode)
14
15  NSUP = point['NSUP']
16  NREL = point['NREL']
17  npcols     = int(nproc / nprows)
18  nproc      = int(nprows * npcols)
19  params = [matrix, 'COLPERM', COLPERM, 'LOOKAHEAD', LOOKAHEAD, 'nthreads',
      nthreads, 'nprows', nprows, 'npcols', npcols, 'NSUP', NSUP, 'NREL', NREL]
20
21  """ pass some parameters through environment variables """
22  info = MPI.Info.Create()
23  envstr= 'OMP_NUM_THREADS=%d\n' %(nthreads)
24  envstr+= 'NREL=%d\n' %(NREL)
25  envstr+= 'NSUP=%d\n' %(NSUP)
26  info.Set('env',envstr)
27  info.Set('npernode','%d'%(npernode))  # YL: npernode is deprecated in openmpi
      4.0, but no other parameter (e.g. 'map-by') works
28
29  """ use MPI spawn to call the executable, and pass the other parameters and
      inputs through command line """
30  comm = MPI.COMM_SELF.Spawn("%s/pddrive_spawn"%(RUNDIR), args=['-c', '%s'%(npcols
      ), '-r', '%s'%(nprows), '-l', '%s'%(LOOKAHEAD), '-p', '%s'%(COLPERM), '%s/%s'%(
      INPUTDIR,matrix)], maxprocs=nproc,info=info)
31
32  """ gather the return value using the inter-communicator, also refer to the
      INPUTDIR/pddrive_spawn.c to see how the return value are communicated """
33  tmpdata = array('f', [0,0])
34  comm.Reduce(sendbuf=None, recvbuf=[tmpdata,MPI.FLOAT],op=MPI.MAX,root=mpi4py.MPI
      .ROOT)
35  comm.Disconnect()
36  retval = tmpdata[0]
37  print(params, ' superlu time: ', retval)
38  return [retval]
39 def main():
40  ntask = 2   # 2 tasks used for MLA
41  NS = 20  # 20 samples per task
42  matrices = ["big.rua", "g4.rua", "g20.rua"]
43  # read the meta.json file
44  (machine, processor, nodes, cores) = GetMachineConfiguration()
45  print ("machine: " + machine + " processor: " + processor + " num_nodes: " + str
      (nodes) + " num_cores: " + str(cores))
46
47  """ Define and print the spaces and constraints """
48  # Task Parameters
49  matrix     = Categoricalnorm (matrices, transform="onehot", name="matrix")
```

```python
50   IS = Space([matrix])
51   # Tuning Parameters
52   COLPERM   = Categoricalnorm (['2', '4'], transform="onehot", name="COLPERM")
53   LOOKAHEAD = Integer      (5, 20, transform="normalize", name="LOOKAHEAD")
54   nprows    = Integer      (1, nprocmax, transform="normalize", name="nprows")
55   npernode  = Integer      (0, int(log2(cores)), transform="normalize", name="
       npernode")
56   NSUP      = Integer      (30, 300, transform="normalize", name="NSUP")
57   NREL      = Integer      (10, 40, transform="normalize", name="NREL")
58   PS = Space([COLPERM, LOOKAHEAD, npernode, nprows, NSUP, NREL])
59   # Output
60   time   = Real           (float("-Inf") , float("Inf"), transform="normalize", name=
       "time")
61   OS = Space([time])
62   # Constraints
63   cst1 = "NSUP >= NREL"
64   cst2 = "nodes * 2**npernode >= nprows"
65   constraints = {"cst1" : cst1, "cst2" : cst2}
66   constants={"nodes":nodes,"cores":cores}
67   print(IS, PS, OS, constraints)
68   problem = TuningProblem(IS, PS, OS, objectives, constraints, None, constants)
69   computer = Computer(nodes = nodes, cores = cores, hosts = None)
70
71   """ Set and validate options """
72   options = Options()
73   options['model_restarts'] = 4   # number of GP models being built in one
       iteration (only the best model is retained)
74   options['distributed_memory_parallelism'] = True # True: Use MPI. One MPI per
       model start in the modeling phase, one MPI per task in the search phase
75   options['shared_memory_parallelism'] = False # True: Use threads. One thread per
        model start in the modeling phase, one MPI per task in the search phase
76   options.validate(computer = computer)
77
78   """ Intialize the tuner with existing data"""
79   data = Data(problem)  # intialize with empty data, but can also load data from
       previous runs
80   gt = GPTune(problem, computer = computer, data = data, options = options,
       driverabspath=os.path.abspath(__file__))
81
82   """ Build MLA with the given list of tasks """
83   giventask = [["g4.rua"], ["g20.rua"]]
84   NI = len(giventask)
85   (data, model,stats) = gt.MLA(NS=NS, NI=NI, Igiven =giventask, NS1 = max(NS//2,1)
       )
86   print("stats: ",stats)
87
88   """ Print all task input and parameter samples """
89   for tid in range(NI):
90     print("tid: %d"%(tid))
91     print("    matrix:%s"%(data.I[tid][0]))
92     print("    Ps ", data.P[tid])
93     print("    Os ", data.O[tid].tolist())
94     print('    Popt ', data.P[tid][np.argmin(data.O[tid])], 'Oopt ', min(data.O[
       tid])[0], 'nth ', np.argmin(data.O[tid]))
95
96   """ Call TLA for a new task using the constructed LCM model"""
```

```
97   newtask = [["big.rua"]]
98   (aprxopts,objval,stats) = gptune.TLA1(newtask, NS=None)
99   print("stats: ",stats)
100
101  """ Print the optimal parameters and function evaluations"""
102  for tid in range(len(newtask)):
103    print("new task: %s"%(newtask[tid]))
104    print('    predicted Popt: ', aprxopts[tid], ' objval: ',objval[tid])
105
106 if __name__ == "__main__":
107   main()
```

Listing 10: superlu_MLA.py.

```
1  int main(int argc, char *argv[])
2  {
3      int       nprow, npcol,lookahead,colperm;
4      char      **cpp, c;
5      FILE *fp;
6    MPI_Comm parent;
7
8      /* Intialize MPI and get the inter communicator. */
9      MPI_Init( &argc, &argv );
10   MPI_Comm_get_parent(&parent);
11
12     /* Read the input and parameters from command line arguments. */
13   for (cpp = argv+1; *cpp; ++cpp) {
14     if ( **cpp == '-' ) {
15       c = *(*cpp+1);
16       ++cpp;
17       switch (c) {
18         case 'h':
19         printf("Options:\n");
20         printf("\t-r <int>: process rows    (default %4d)\n", nprow);
21         printf("\t-c <int>: process columns (default %4d)\n", npcol);
22         exit(0);
23         break;
24         case 'r': nprow = atoi(*cpp);     // number of row processes
25           break;
26         case 'c': npcol = atoi(*cpp);     // number of column processes
27           break;
28         case 'l': lookahead = atoi(*cpp); // size of lookahead window
29           break;
30         case 'p': colperm = atoi(*cpp);   // column permutation
31           break;
32       }
33     } else { /* Last arg is considered a filename */
34       if ( !(fp = fopen(*cpp, "r")) ) {          // the file storing the sparse
     matrix
35           ABORT("File does not exist");
36       }
37       break;
38     }
39     }
40
41   /* Read the input and parameters from environment variables (including NSUP,
```

```
     NREL and OMP_NUM_THREADS) */
42  if (master process) {
43    print_sp_ienv_dist(&options);
44    print_options_dist(&options);
45    fflush(stdout);
46  }
47
48  /* Allocate superlu meta-data and call the computation routine. */
49  //...
50
51  /* sending the results (numerical factorization time) to the parent process */
52  result = runtime results;
53  MPI_Reduce(result, MPI_BOTTOM, 1, MPI_FLOAT,MPI_MAX, 0, parent);
54
55  /* DEALLOCATE SupreLU meta-data. */
56  //...
57
58  /* Disconnect the inter communicator and finalize the intra communicator. */
59  MPI_Comm_disconnect(&parent);
60    MPI_Finalize();
61 }
```

Listing 11: pddrive_spawn.c.

5.2.3 Muti-objective MLA

The following example demonstrates the capability of multi-objective auto-tuning feature of GP-Tune with two objectives (runtime and memory of a sparse LU factorization). The example calls MLA to build a LCM model per objective for the SuperLU_DIST driver pddrive_spawn.c using three tasks [["big.rua"], ["g4.rua"], ["g20.rua"]]. Note that only the simplified code is shown here, please refer to GPTune/examples/SuperLU_DIST/superlu_MLA_MO.py and GPTune/examples/-SuperLU_DIST/superlu_dist/EXAMPLE/pddrive_spawn.c for the complete working codes.

```
1  def objectives(point):
2    nodes = point['nodes']
3    cores = point['cores']
4    RUNDIR = os.path.abspath(__file__ + "/../superlu_dist/build/EXAMPLE") # the path
         to the executable
5    INPUTDIR = os.path.abspath(__file__ + "/../superlu_dist/EXAMPLE/") # the path to
         the matrix collection
6    matrix = point['matrix']
7    COLPERM = point['COLPERM']
8    LOOKAHEAD = point['LOOKAHEAD']
9    nprows = point['nprows']
10   npernode = 2**point['npernode']
11   nproc = nodes*npernode
12   nthreads = int(cores / npernode)
13   NSUP = point['NSUP']
14   NREL = point['NREL']
15   npcols     = int(nproc / nprows)
16   nproc     = int(nprows * npcols)
17   params = [matrix, 'COLPERM', COLPERM, 'LOOKAHEAD', LOOKAHEAD, 'nthreads',
       nthreads, 'nprows', nprows, 'npcols', npcols, 'NSUP', NSUP, 'NREL', NREL]
18
19
```

```python
     """ pass some parameters through environment variables """
20   info = MPI.Info.Create ()
21   envstr= 'OMP_NUM_THREADS=%d\n' %(nthreads)
22   envstr+= 'NREL=%d\n' %(NREL)
23   envstr+= 'NSUP=%d\n' %(NSUP)
24   info.Set('env',envstr)
25   info.Set('npernode','%d'%(npernode))  # YL: npernode is deprecated in openmpi
26      4.0, but no other parameter (e.g. 'map-by') works
27
28   """ use MPI spawn to call the executable, and pass the other parameters and
        inputs through command line """
29   comm = MPI.COMM_SELF.Spawn("%s/pddrive_spawn"%(RUNDIR), args=['-c', '%s'%(npcols
        ), '-r', '%s'%(nprows), '-l', '%s'%(LOOKAHEAD), '-p', '%s'%(COLPERM), '%s/%s'%(
        INPUTDIR,matrix)], maxprocs=nproc,info=info)
30
31   """ gather the return value using the inter-communicator, also refer to the
        INPUTDIR/pddrive_spawn.c to see how the return value are communicated """
32   tmpdata = array('f', [0,0])
33   comm.Reduce(sendbuf=None, recvbuf=[tmpdata,MPI.FLOAT],op=MPI.MAX,root=mpi4py.MPI
        .ROOT)
34   comm.Disconnect ()
35
36   print(params, ' superlu time: ', tmpdata[0], ' memory: ', tmpdata[1])
37   return tmpdata
38
39
40
41 def main():
42   ntask = 3   # 3 tasks used for MLA
43   NS = 20  # 20 samples per task
44   matrices = ["big.rua", "g4.rua", "g20.rua"]
45   (machine, processor, nodes, cores) = GetMachineConfiguration ()
46   print ("machine: " + machine + " processor: " + processor + " num_nodes: " + str
        (nodes) + " num_cores: " + str(cores))
47
48
49   """ Define and print the spaces and constraints """
50   # Task Parameters
51   matrix    = Categoricalnorm (matrices, transform="onehot", name="matrix")
52   IS = Space([matrix])
53   # Tuning Parameters
54   COLPERM   = Categoricalnorm (['2', '4'], transform="onehot", name="COLPERM")
55   LOOKAHEAD = Integer     (5, 20, transform="normalize", name="LOOKAHEAD")
56   nprows    = Integer     (1, nprocmax, transform="normalize", name="nprows")
57   npernode  = Integer     (0, int(log2(cores)), transform="normalize", name="
        npernode")
58   NSUP      = Integer     (30, 300, transform="normalize", name="NSUP")
59   NREL      = Integer     (10, 40, transform="normalize", name="NREL")
60   PS = Space([COLPERM, LOOKAHEAD, npernode, nprows, NSUP, NREL])
61   # Output
62   runtime   = Real        (float("-Inf") , float("Inf"), transform="normalize",
        name="r")
63   memory    = Real        (float("-Inf") , float("Inf"), transform="normalize",
        name="memory")
64   OS = Space([runtime, memory])
```

```
65    # Constraints
66    cst1 = "NSUP >= NREL"
67    cst2 = "nodes * 2**npernode >= nprows"
68    constraints = {"cst1" : cst1, "cst2" : cst2}
69    constants ={"nodes":nodes,"cores":cores}
70    print(IS, PS, OS, constraints)
71    problem = TuningProblem(IS, PS, OS, objectives, constraints, None, constants)
72    computer = Computer(nodes = nodes, cores = cores, hosts = None)
73
74    """ Set and validate options """
75    options = Options()
76    options['model_restarts'] = 1   # number of GP models being built in one
        iteration (only the best model is retained)
77    options['distributed_memory_parallelism'] = False # True: Use MPI. One MPI per
        model start in the modeling phase, one MPI per task in the search phase
78    options['shared_memory_parallelism'] = False # True: Use threads. One thread per
        model start in the modeling phase, one MPI per task in the search phase
79    options['search_algo'] = 'nsga2' # multi-objective search algorithm
80    options['search_pop_size'] = 1000 # Population size in pgymo
81    options['search_gen'] = 10 # Number of evolution generations in pgymo
82    options['search_best_N'] = 4 # Maximum number of points selected using a multi-
        objective search
83    options.validate(computer = computer)
84
85    """ Intialize the tuner with existing data"""
86    data = Data(problem) # intialize with empty data, but can also load data from
        previous runs
87    gt = GPTune(problem, computer = computer, data = data, options = options)
88
89    """ Building MLA with the given list of tasks """
90    giventask = [["big.rua"], ["g4.rua"], ["g20.rua"]]
91    NI = len(giventask)
92    (data, models,stats) = gt.MLA(NS=NS, NI=NI, Igiven =giventask, NS1 = max(NS
        //2,1))
93    print("stats: ",stats)
94
95    """ Print all task input and parameter samples; search for and print the Pareto
        front"""
96    for tid in range(NI):
97      print("tid: %d"%(tid))
98      print("    matrix:%s"%(data.I[tid][0]))
99      print("    Ps ", data.P[tid])
100     print("    Os ", data.O[tid])
101     ndf, dl, dc, ndr = pg.fast_non_dominated_sorting(data.O[tid])
102     front = ndf[0]
103     # print('front id: ',front)
104     fopts = data.O[tid][front]
105     xopts = [data.P[tid][i] for i in front]
106     print('    Popts ', xopts)
107     print('    Oopts ', fopts)
108
109 if __name__ == "__main__":
110   main()
```

Listing 12: superlu_MLA_MO.py.

51

5.3 SuperLU_DIST (RCI)

For the RCI mode, in addition to the GPTune Python driver, the user also needs to provide a GPTune bash driver. The bash driver will keep querying the GPTune Python driver for next sampling points, get the sampling points for the dababase file, invoke the application, and write results back into the dabasefile. This process continues until there is no more sample required. Following the database format described in Section 3.3.2, "time: null, memory: null" will appear under "evaluation_result: ", indicating this is a sample requiring evaluation. As the application is invoked from bash, it doesn't need to be compiled using the same software dependence as GPTune, hence OpenMPI is not required. In addition, the modification in Section 4.1 is also not needed. For the SuperLU_DIST application, all tuning and task parameters can be passed to the application via environment variables and command line options, and the objective function evaluations can be obtained by searching the runlog.

5.3.1 Preparing the meta JSON file

Just like *PDGEQRF*, a meta JSON file located at ./gptune/meta.json needs to be edited for SuperLU_DIST.

5.3.2 Muti-objective MLA in RCI mode

The following example performs the same tuning experiment as Section 5.2.3. In other words, the example demonstrates the capability of multi-objective auto-tuning feature of GPTune (in the RCI mode) with two objectives (runtime and memory of a sparse LU factorization). The example calls MLA to build a LCM model per objective for the SuperLU_DIST driver pddrive_spawn.c using three tasks [["big.rua"], ["g4.rua"], ["g20.rua"]]. Note the bash scripts keeps calling superlu_MLA_MO_RCI.py which is similar to the script superlu_MLA_MO.py in Section 5.2, except that superlu_MLA_MO_RCI.py doesn't define the objective function (as it's directly executed in the bash script), and options['RCI_mode']=True is required. Note that only the simplified code is shown here, please refer to GPTune/examples/SuperLU_DIST_RCI/superlu_MLA_MO_RCI.py and GPTune/examples/SuperLU_DIST_RCI/ superlu_MLA_MO_RCI.sh.

```
1  start=`date +%s`
2  nrun=20 # number of samples per task
3  # name of your machine, processor model, number of compute nodes, number of cores
        per compute node, which are defined in .gptune/meta.json
4  declare -a machine_info=($(python -c "from gptune import *;
5  (machine, processor, nodes, cores)=list(GetMachineConfiguration());
6  print(machine, processor, nodes, cores)"))
7  machine=${machine_info[0]}
8  processor=${machine_info[1]}
9  nodes=${machine_info[2]}
10 cores=${machine_info[3]}
11
12 obj1=time     # name of the first objective defined in the python file
13 obj2=memory    # name of the second objective defined in the python file
14
15 database="gptune.db/SuperLU_DIST.json"  # the phrase SuperLU_DIST should match the
        application name defined in .gptune/meta.jason
16
17 # start the main loop
```

```
18  more=1
19  while [ $more -eq 1 ]; do
20
21      # call GPTune and ask for next sample points
22      python ./superlu_MLA_MO_RCI.py -nrun $nrun
23
24      # check whether GPTune needs more data
25      idx=$( jq -r --arg v0 $obj1 '.func_eval | map(.evaluation_result[$v0] == null)
    | index(true) ' $database )
26      if [ $idx = null ];then; more=0; fi;
27
28      # if so, call the application code (GPTune can requires mutiple samples to be
    evaluated)
29      while [ ! $idx = null ]; do
30          echo " $idx"     # idx indexes the record that has null objective function
    values
31
32          # use jq to retrieve task and tuning parameters
33          declare -a input_para=($( jq -r --argjson v1 $idx '.func_eval[$v1].
    task_parameter' $database | jq -r '.[]'))
34          declare -a tuning_para=($( jq -r --argjson v1 $idx '.func_eval[$v1].
    tuning_parameter' $database | jq -r '.[]'))
35
36          # get the task input parameters, the parameters should follow the sequence
     of definition in superlu_MLA_MO_RCI.py
37          matrix=${input_para[0]}
38
39          # get the tuning parameters, the parameters should follow the sequence of
    definition in superlu_MLA_MO_RCI.py
40          COLPERM=${tuning_para[0]}
41          LOOKAHEAD=${tuning_para[1]}
42          npernode=${tuning_para[2]}
43          nprows=${tuning_para[3]}
44          NSUP=${tuning_para[4]}
45          NREL=${tuning_para[5]}
46
47          # set environment variables and command line options
48          npernode=$((2**$npernode))
49          export OMP_NUM_THREADS=$(($cores / $npernode))
50          export NREL=$NREL
51          export NSUP=$NSUP
52          nproc=$(($nodes*$npernode))
53          npcols=$(($nproc / $nprows))
54          RUNDIR="../SuperLU_DIST/superlu_dist/build/EXAMPLE"
55          INPUTDIR="../SuperLU_DIST/superlu_dist/EXAMPLE/"
56
57          # run the application, this doesn't have to be openmpi-compiled code
58          mpirun -n $nproc $RUNDIR/pddrive_spawn -c $npcols -r $nprows -l $LOOKAHEAD
    -p $COLPERM $INPUTDIR/$matrix | tee a.out
59
60          # get the result (for this example: search the runlog)
61          result1=$(grep 'Factor time' a.out | grep -Eo '[+-]?[0-9]+([.][0-9]+)?')
62          result2=$(grep 'Total MEM' a.out | grep -Eo '[+-]?[0-9]+([.][0-9]+)?')
63
64          # use jq to write the data back to the database file
```

```
65          jq --arg v0 $obj1 --argjson v1 $idx --argjson v2 $result1 '.func_eval[$v1].
     evaluation_result[$v0]=$v2' $database > tmp.json && mv tmp.json $database
66          jq --arg v0 $obj2 --argjson v1 $idx --argjson v2 $result2 '.func_eval[$v1].
     evaluation_result[$v0]=$v2' $database > tmp.json && mv tmp.json $database
67
68          # get the next sample to be evaluated
69          idx=$( jq -r --arg v0 $obj1 '.func_eval | map(.evaluation_result[$v0] ==
     null) | index(true) ' $database )
70       done
71 done
72 end=`date +%s`
73 runtime=$((end-start))
74 echo "Total tuning time: $runtime"
```

Listing 13: superlu_MLA_MO_RCI.sh.

6   Numerical experiments

6.1   Parallel speedups of GPTune

Consider the following model problem to be tuned, with the objective function given explicitly as

$$y_{demo}(t, x) = \exp\left(-(x+1)^{t+1}\right)\cos(2\pi x)\sum_{i=1}^{3}\sin\left(2\pi x(t+2)^{i}\right) \tag{7}$$

where t and x denote the task and tuning parameters. Note that this function is highly non-convex and we are interested in finding the minimum for $x \in [0, 1]$ for multiple tasks t. Fig. 4 plots $y_{demo}(t, x)$ versus x for four different values of t and marks the minimum objective function values.

First, we evaluate the parallel performance of the MLA algorithm on a Threadripper 1950X 16-core processor using $\delta = 20$ tasks. In Fig. 5, we plot the runtime of the modeling and search phases using 1 and 16 cores, by enabling the GPTune parameter distributed_memory_parallelism. For simplicity, we set the initial random sample count to NS1 = NS-1 (i.e., only one MLA iteration is performed). As we increase the number of total samples NS from 10 to 160 (with the LCM kernel matrix size changing from 200 to 3200), 13X (comparing the two blue curves) and 10X (comparing the two black curves) speedups are observed for the modeling and sampling phases, respectively.

6.2   Advantage of using performance models

Next, we evaluate the effects of the performance models using the above objective function $y_{demo}(t, x)$. We test three performance models seperately, $\tilde{y}_1(t, x) = y_{demo}(t, x)$ (the model is excatly the objective), $\tilde{y}_2(t, x) = 10y_{demo}(t, x)$ (the model output is a factor of 10 larger than the objective), and $\tilde{y}_3(t, x) = (1 + 0.1 \times r(x))y_{demo}(t, x)$ (the model is the objective with random scaling factors). Here $r(x)$ is a random number drawn from the normal distribution $\mathcal{N}(0, 1)$. We set NS1 = NS/2 and use a single task $t = 6$. Table 2 lists the minimum objective value returned by GPTune's MLA algorithm with varying total sample counts NS. Without any performance model, MLA has still not found the minimum after 640 samples. With the exact model $\tilde{y}_1$ (which is not practical), not surprisingly, at most 20 samples are sufficient to get very close to the minimum: -4.89E-01 (see Fig. 4(d)). With the scaled models $\tilde{y}_2$ and $\tilde{y}_3$, at most 80 samples and 40 samples are sufficient,

54

Figure 4: The objective functions in (7) for four task parameter values t.

Figure 5: Modeling and search time for 1 and 16 MPIs using the objective function f_{demo}.

respectively. In other words, with the coarse performance models, GPTune requires significantly fewer samples to build an accurate LCM model.

NS	10	20	40	80	160	320	640
None	-1.40E-01	-6.06E-02	-2.93E-02	-3.79E-01	-2.69E-01	-4.25E-01	-3.83E-01
$\tilde{y}_1$	-4.51E-01	-4.88E-01	-4.85E-01	-4.88E-01	-4.86E-01	-4.89E-01	-4.89E-01
$\tilde{y}_2$	-2.98E-01	-3.72E-01	-4.83E-01	-4.89E-01	-4.89E-01	-4.88E-01	-4.89E-01
$\tilde{y}_3$	-4.52E-01	-4.52E-01	4.88E-01	-4.77E-01	-4.89E-01	-4.89E-01	-4.89E-01

Table 2: Minimum found by GPTune for the objective f_{demo} with and without performance models.

6.3 Efficiency of multi-task learning

Next, we use the ScaLAPACK QR example in Section 5.1.2 to compare the performance of the GPTune MLA algorithms with single-task (δ=1) and multi-task (δ=20) settings. We use 16 NERSC Cori nodes assuming a fixed budget of $\delta \times$ NS = 400 and NS1 = NS/2. For δ=1, we consider the task ($m = 4674, n = 3608$); for δ=20, we also consider 19 other tasks that are randomly generated with $m, n < 5000$ (in practice, one may choose all 20 tasks of interest).

Table 3 shows the runtime breakdown of the single-task and multi-task MLA algorithms. For this example, the total runtime is dominated by the objective function evaluation. The multi-task MLA requires less objective evaluation time as it involves 19 other less expensive tasks. In addition, the multi-task modeling phase is much faster than single-task one as it requires fewer MLA iterations. More specifically, the multi-task modeling requires 10 iterations with the LCM matrix dimensions $200, 220, 240, ..., 380$ while the single-task modeling requires 200 iterations with the LCM matrix dimensions $200, 201, 202, ..., 399$.

Fig. 6 plots the runtime (obtained through runing the application) and corresponding GFlops using the optimal tuning parameters for all 20 tasks. The red dots correspond to δ=20 and the blue dots correspond to δ=1. Note that the surfaces are constructed via the Matlab "griddata" function using the red dots. The multi-task MLA not only achieves a very similar minimum to the single-task MLA for ($m = 4674, n = 3608$), but also finds minima for all the other 19 tasks.

6.4 Capability of multi-objective tuning

Finally, we illustrate the multi-objective feature of GPTune for tuning the factorization time and memory in SuperLU_DIST [16]. As explained in Section 5.2, we consider six tuning parameters (COLPERM, LOOKAHEAD, nproc, nprows, NSUP, NREL) and two objectives (time, memory). As an example, we apply sparse factorization to a matrix "Si2" (single task) from the SuiteSparse Matrix Collection [4] using 8 NERSC Cori nodes.

	total time	objective evaluation	modeling	search
Single-task	15092.4	14062.3	907.8	120.1
Multi-task	9386.8	9091.4	85.7	208.1

Table 3: Runtime of different phases in the GPTune single-objective and multi-objective MLA with a total of 400 samples.
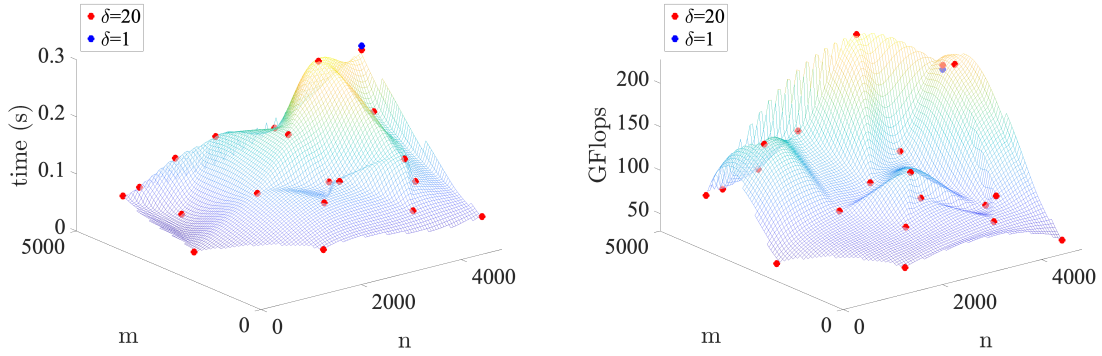
Figure 6: (a) Runtime (the objective) and (b) GFlops for 20 tasks of QR factorization after auto-tuning.

As a reference, we also consider single-objective (time) and (memory). For example, single-objective (memory) tuning means minimizing the memory usage ignoring the impact on runtime (as long as the code still runs correctly). Table 4 lists the default and optimal (single-objective) tuning parameters. The default parameters are those used by SuperLU_DIST without any tuning. The optimal ones are vastly different from the default ones.

Fig. 7 plots the objective function values (via running the application) on the logarithmic scale using the default tuning parameters, and those returned by the GPTune single-objective and multi-objective MLA algorithms. The multi-objective MLA algorithm returns multiple tuning parameter configurations and their objective function values (in black), among which no data point dominates over any other in both objectives. In other words, the black dots lie on the Pareto front. We see that the single-objective minima (in yellow and magenta) lie on or near the Pareto front formed by the multi-objective minima (in black). Not surprisingly, the default objective values (in cyan) are far from optimal in either dimension.

	COLPERM	LOOKAHEAD	nproc	nprows	NSUP	NREL
Default	4	10	256	16	128	20
Single-objective (time)	2	6	216	149	295	37
Single-objective (memory)	2	5	193	20	31	22

Table 4: Default tuning parameters and optimal ones returned by the GPTune single-objective MLA algorithm.

References

[1] mpi4py. `https://pypi.org/project/mpi4py/`.

[2] PyGMO. `https://esa.github.io/pygmo/`.

[3] Scikit-Optimize. `https://scikit-optimize.github.io`.

Figure 7: Logarithmic plots of the optimal objective functions values (factorization time and memory of SuperLU_DIST with 8 NERSC Cori nodes) found by GPTune using single-objective and multi-objective tuning. The objective function values using the default tuning parameters are also plotted.

[4] Suitesparse matrix collection. `https://sparse.tamu.edu/`.

[5] L. S. Blackford, J. Choi, A. Cleary, E. D'Azevedo, J. Demmel, I. Dhillon, J. Dongarra, S. Hammarling, G. Henry, A. Petitet, K. Stanley, D. Walker, and R. C. Whaley. *ScaLAPACK Users' Guide*. Society for Industrial and Applied Mathematics, Philadelphia, PA, 1997.

[6] K. Chodorow. *MongoDB: The Definitive Guide: Powerful and Scalable Data Storage*. O'Reilly Media, 2013.

[7] K. Deb, A. Pratap, S. Agarwal, and T. Meyarivan. A fast and elitist multiobjective genetic algorithm: Nsga-ii. *IEEE transactions on evolutionary computation*, 6(2):182–197, 2002.

[8] R. D. Falgout and U. M. Yang. hypre: A library of high performance preconditioners. In P. M. A. Sloot, A. G. Hoekstra, C. J. K. Tan, and J. J. Dongarra, editors, *Computational Science — ICCS 2002*, pages 632–641, Berlin, Heidelberg, 2002. Springer Berlin Heidelberg.

[9] S. Falkner, A. Klein, and F. Hutter. BOHB: Robust and efficient hyperparameter optimization at scale. In J. Dy and A. Krause, editors, *Proceedings of the 35th International Conference on Machine Learning*, volume 80 of *Proceedings of Machine Learning Research*, pages 1437–1446. PMLR, 10–15 Jul 2018.

[10] P. Frazier. A Tutorial on Bayesian Optimization. `https://arxiv.org/abs/1807.02811`, 2018.

[11] GPy. GPy: A gaussian process framework in python. `http://github.com/SheffieldML/GPy`, since 2012.

[12] Jason Ansel and Shoaib Kamil and Kalyan Veeramachaneni and Jonathan Ragan-Kelley and Jeffrey Bosboom and Una-May O'Reilly and Saman Amarasinghe. OpenTuner: An Extensible Framework for Program Autotuning. In *International Conference on Parallel Architectures and Compilation Techniques*, Edmonton, Canada, August 2014.

[13] D. R. Jones, M. Schonlau, and W. J. Welch. Efficient global optimization of expensive black-box functions. *Journal of Global optimization*, 13(4):455–492, 1998.

[14] json.org. Json. `https://www.json.org/json-en.html`.

[15] J. Kennedy and R. Eberhart. Particle swarm optimization. In *Proceedings of ICNN'95 - International Conference on Neural Networks*, volume 4, pages 1942–1948 vol.4, Nov 1995.

[16] X. S. Li and J. W. Demmel. SuperLU_DIST: a scalable distributed-memory sparse direct solver for unsymmetric linear systems. *ACM Trans. Math. Softw.*, 29(2):110–140, June 2003.

[17] D. C. Liu and J. Nocedal. On the limited memory bfgs method for large scale optimization. *MATHEMATICAL PROGRAMMING*, 45:503–528, 1989.

[18] W. Sid-lakhdar, M. Aznaveh, X. Li, and J. Demmel. Multitask and Transfer Learning for Autotuning Exascale Aplications. `https://arxiv.org/abs/1908.05792`, 2019.

[19] ytopt. ytopt: Machine-learning-based search methods for autotuning. `https://github.com/ytopt-team/ytopt`, 2019.